

HEII REIONIZATION AND ITS EFFECT ON THE IGM

MATTHEW MCQUINN¹, ADAM LIDZ¹, MATIAS ZALDARRIAGA^{1,2}, LARS HERNQUIST¹, PHILIP F. HOPKINS¹, SUVENDRA DUTTA¹, CLAUDE-ANDRÉ FAUCHER-GIGUÈRE¹*Draft version December 30, 2008*

ABSTRACT

Observations of the intergalactic medium (IGM) suggest that quasars reionize HeII in the IGM at $z \approx 3$. We have run a set of 190 and 430 comoving Mpc simulations of HeII being reionized by quasars to develop an understanding of the nature of HeII reionization and its potential impact on observables. We find that HeII reionization heats regions in the IGM by as much as 25,000 K above the temperature that is expected otherwise, with the volume-averaged temperature increasing by $\sim 12,000$ K and with large temperature fluctuations on ~ 50 Mpc scales. Much of this heating occurs far from quasars by photons with long mean free paths. We find a temperature-density equation of state of $\gamma - 1 \approx 0.3$ during HeII reionization, but with a wide dispersion in this relation having $\sigma_T \sim 10^4$ K. HeII reionization by the observed population of quasars cannot produce an inverted relation ($\gamma - 1 < 0$). Our simulations are consistent with the observed evolution in the mean transmission of the HeII Ly α forest. We argue that the heat input from HeII reionization is unable to cause the observed depression at $z \approx 3.2$ in the HI Ly α forest opacity as has been suggested. We investigate how uncertainties in the properties of QSOs and of HeII Lyman-limit systems influence our predictions.

Subject headings: cosmology: theory – intergalactic medium

1. INTRODUCTION

In the standard picture for the reionization history of the Universe, radiation from Population II stars ionized the intergalactic HI at $z > 6$ as well as the HeI, converting the vast majority of the intergalactic helium to HeII. However, these stars cannot ionize HeII, and at $z \approx 3$ quasars, with their harder UV spectrum, doubly ionize the intergalactic helium. To test this model, many observations are targeting high redshifts to probe hydrogen reionization (e.g., Fan et al. 2002; Taniguchi et al. 2005; Bouwens et al. 2007; Kashikawa et al. 2006; Stark et al. 2007; Totani et al. 2006). In this picture, HeII reionization occurs at redshifts for which there is substantially more data on the state of the intergalactic medium (IGM).

In fact, a number of observations suggest that HeII reionization happened at $z \sim 3$. Two measurements of the mean transmission in the HI Ly α forest have noted an upward bump at $z \approx 3.2$ (Bernardi et al. 2003; Faucher-Giguère et al. 2008a), which Theuns et al. (2002b) interpreted as arising from a temperature increase of the IGM during HeII reionization (but see Faucher-Giguère et al. (2008a) for alternative explanations). An increase in the average temperature of the IGM would also decrease the small-scale fluctuations in the HI Ly α forest. Ricotti et al. (2000) and Schaye et al. (2000) measured the temperature from the widths of the narrowest lines in the HI Ly α forest and claimed to have detected a sudden increase in the temperature of $\Delta T \sim 10^4$ K between $z = 3.5$ and 3. Photo-heating during HeII reionization is the only known process that could be responsible for such an effect (e.g., Miralda-Escudé & Rees 1994; Abel & Haehnelt 1999).

However, a subsequent study by McDonald et al. (2001) using a similar method and Zaldarriaga et al. (2001) using the HI Ly α forest power spectrum did not confirm this sudden increase in temperature, but rather found a constant temperature at mean density of $T_0 \approx 17,000$ K for $2 < z < 4$. Temperatures of 17,000 K are difficult to explain without HeII reionization occurring at $z \sim 3$ (Hui & Haiman 2003), and it is unclear whether as sudden an increase in temperature as Ricotti et al. (2000) and Schaye et al. (2000) find is even expected theoretically.

If a substantial fraction of the helium is in HeII ($\gtrsim 1\%$), this would produce a Gunn-Peterson absorption trough in the spectra of high-redshift quasars at wavelengths blueward of HeII Ly α . Observations of HeII Ly α forest absorption at $2.8 < z < 3.3$ find 10s of comoving Mpc regions with no detected transmission (Jakobsen et al. 1994; Davidsen et al. 1996; Hogan et al. 1997; Reimers et al. 1997; Heap et al. 2000). These troughs may signify the presence of diffuse intergalactic HeII. However, current data, which consist of only a few quasar sight-lines, do not rule out the intergalactic HeII being primarily ionized and in photo-ionization equilibrium with a weak background (Giroux & Shull 1997; Fardal et al. 1998; Heap et al. 2000). The Cosmic Origins Spectrograph, which NASA plans to install on the Hubble Space Telescope in 2009, will increase the quantity and quality of HeII Ly α forest sight-lines.

As the intergalactic HeII becomes progressively more ionized, the extragalactic UV background will harden around the ionization energy of HeII at 54.4 eV. This hardening will affect the ionization state of intergalactic metals. Songaila (1998) observed a sharp evolution at $z \approx 3$ in the column density ratios in SIV (Ionization Potential = 45.1 eV) to CIV (64.5 eV) absorbers. Boksenberg et al. (2003) found evidence for a more gradual hardening of the background between $2 < z < 4$ from the column density ratios of NV (98 eV) to

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138; mmcquinn@cfa.harvard.edu

² Jefferson Laboratory of Physics, Harvard University, Cambridge, MA 02138

CIV. Finally, by simultaneously fitting to multiple metal lines that originate from the same absorption systems, Agafonova et al. (2005) and Agafonova et al. (2007) inferred a background spectrum that is hardening at $z \approx 3$ near 4 Ry.

Measurements from the $z \sim 3$ HI Ly α forest ignore the effects of a patchy HeII reionization process. For example, estimates of the photo-ionizing background and the IGM temperature from the forest assume a power-law temperature-density relation. Cosmological parameter studies from the HI Ly α forest power spectrum make a similar assumption. Different regions can have vastly different HeII reionization histories, resulting in a more complicated distribution of temperatures and pressure-smoothing scales than is commonly adopted. Realistic simulations of HeII reionization will help quantify the level at which this process biases these measurements.³

The aim of this paper is to run realistic simulations of HeII reionization to understand the morphology of this process as well as its effect on observables. We concentrate primarily on its impact on quantities that are sensitive to the IGM temperature, but we also study the effect of HeII reionization on the transmission in the HeII Ly α forest.

Sokasian et al. (2002) and Paschos et al. (2007) have performed the most realistic simulations of patchy HeII reionization to date. There are several differences between our work and these earlier investigations. Both of these studies employed volumes $\leq 100^3$ comoving Mpc³. Here, we examine HeII reionization in 186³ and 429³ comoving Mpc³ volumes, providing a more representative cosmic sample. However, both Sokasian et al. (2002) and Paschos et al. (2007) simulated HeII reionization as a post-processing step on top of cosmological simulations that included gas dynamics. Our study instead uses N-body simulations, which result in a less realistic model for the gas distribution, but afford a larger dynamic range. Furthermore, Sokasian et al. (2002) assumed sharp ionizing fronts and ignored the detailed temperature evolution. Paschos et al. (2007) did not calculate the gas temperature self-consistently. Our calculations capture the width of the ionization fronts and the temperature in a consistent manner. Finally, in contrast to previous studies, our work presents a large set of radiative transfer simulations in order to survey the parameter space.

In Section 2, we describe the details of our code. The models for the quasar sources are described in Section 3. Section 4 presents the simulations. Finally, Section 5 addresses the implications HeII reionization has on observations of the HI and HeII Ly α forests.

Throughout, we use a Λ CDM cosmology with $n_s = 1$, $\sigma_8 = 0.8$, $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$, $\Omega_b = 0.046$, and $h = 0.7$, consistent with the most recent WMAP results (Komatsu et al. 2008). All distances are in comoving units unless specified otherwise. An overbar over a variable signifies a volume average, and x_Y is the fraction of helium/hydrogen that is in ionization state Y .

2. ALGORITHM

³ Of note, Lai et al. (2006) used simple models for HeII reionization to show that it has a surprisingly small effect on the Ly α forest power spectrum on large scales, modifying it at a $\lesssim 5\%$ level for wavevectors $k \lesssim 5$ comoving Mpc⁻¹.

Our work employs a new ray-tracing code that has been adapted significantly from the one originally presented in Sokasian et al. (2001) and refined in McQuinn et al. (2007). This code performs cosmological radiative transfer as a post-processing step on top of N-body or Smooth Particle Hydrodynamics density fields. In this study we use N-body simulations. The differences between our code and those described in Sokasian et al. (2001) and McQuinn et al. (2007) include that it tracks photons in multiple frequency bins, that it calculates the temperature of the gas, that it does not assume sharp ionizing fronts, and that it is parallelized over shared memory. This section describes the details of our code, and Appendix B discusses various tests of it.

At the beginning of each timestep, the code grids the particles from an N-body snapshot, and inputs a list of source positions and luminosities. Then, the code adjusts the neutral fraction to account for the total number of recombinations that will occur over the ensuing timestep. Next, the code casts rays from every source, randomizing the order of the sources in this loop as well as the direction in which the rays are cast. Every ray carries a set number of photons. Initially, $N_{\text{ray}} = 12 \times 4^L$ rays are cast with $L = 5$ or 6 for an isotropically emitting source, with these rays uniformly pixelating the unit sphere using the HealPIX algorithm as described in Abel & Wandelt (2002).⁴ Each source must wait for all of the other sources to send a ray before it sends a second ray. Rays are split adaptively as they traverse the box. At a minimum, one ray from a source intersects every cell face for cells within the light travel time of a ray. (Although cells that are tens of Mpc from a quasar typically receive multiple rays from the source in a timestep, owing to the large numbers of rays that are initially sent.)

A ray travels either until its photons are expended⁵ or until it has traveled a distance $c\Delta t$, where Δt is the simulation timestep. Rays that have traveled $c\Delta t$ are stored onto disk until the following timestep, at which time the stored rays are redshifted in frequency, randomly ordered, and recast. Rays that have traveled a total of 1.5 box lengths are terminated and their photons are added to the background (Section 2.1). Most cosmological radiative transfer codes allow rays to travel further within a timestep than $c\Delta t$. To study HeII reionization, it is crucial to capture light-travel effects.

For the calculations presented here, rays carry photons with energies between the ionization potential of HeII ($E_{\text{HeII}} = 54.4$ eV) and the energy that has a mean free path (m.f.p.) equal to 3.5 times the size of the 186 Mpc simulation box (1.8 times for the 429 Mpc box), which works out to $E_\gamma \sim 500$ eV in the 186 Mpc box at $z = 3$.

⁴ While the value for N_{ray} may seem large, most of the computation time is spent tracing rays that are far from the source so that the speed of the algorithm depends weakly on the number of rays that are initially cast.

⁵ The termination criterion is that the number of photons in the ray be ϵ times the number of atoms in a fully neutral cell with baryonic overdensity equal to zero, where we set $\epsilon \sim 10^{-4} - 10^{-5}$ with the exact value depending on the run in question. The photons from terminated rays are placed in the background. We have run a simulation with $\epsilon = 10^{-6}$ to test convergence (simulation L1d in Table 1) and find no notable differences with simulation L1 for which $\epsilon = 8 \times 10^{-4}$. In fact, the temperature and ionization fields are almost visually indistinguishable between these two simulations.

The number 3.5, while somewhat arbitrary, assures that the flux is uniform on the box scale for photon energies that are not tracked by the rays. Photons up to 5 keV are put into an ionizing background (§2.1). Many of the background photons are never absorbed because photons with $E_\gamma \approx 0.8 \times [(1+z)/4]^{1/2} \bar{x}_{\text{HeII}}^{1/3}$ keV have a m.f.p. equal to the Hubble scale.

Each ray carries some number of photons with energy at E_i for $1 \leq i \leq n_\nu$. This study typically uses $n_\nu = 5$. The value of E_γ of the photons within an energy bin is specified to conserve photon number and total energy given the source spectrum. The centers of the energy bins that are tracked by the rays are roughly (60, 70, 100, 140, 270) eV for our fiducial spectrum in the 186 Mpc box. In Appendix B, we show that 5 is the minimum number of frequency bins that provides adequate convergence. The final volume-averaged temperature of a simulation with a much larger number of energy bins is within 300 K of the same run with $n_\nu = 5$ energy bins, and the ionization and temperature fields are similar.

Ideally, radiative transfer codes would employ a timestep that is much smaller than the photo-ionization equilibrium timescale. For cosmological ray tracing codes, such small timesteps can be computationally intractable (although, see Trac & Cen 2007). Our algorithm is designed to operate with a timestep that can be much larger than the equilibration timescale; although, in this study it is typically used in the regime where the two are comparable. The timestep used in our simulations is $\Delta t \approx 10$ Myr. Compare this to the photo-ionization equilibration timescale $(\Gamma_{\text{HeII}} + \alpha n_e)^{-1}$, where Γ_{HeII} is the HeII photo-ionization rate and α is the recombination rate coefficient. This timescale, which in most cases is dominated by the Γ_{HeII} term, is observed to be 3–30 Myr at $2.5 < z < 3$ (i.e. $-15 < \log_{10} \Gamma_{\text{HeII}} < -14$ in c.g.s.), but it can be much shorter near quasars. An accurate solution can be obtained with our time steps as demonstrated in Appendix C.

Converging to the solution also requires capturing the correct order that rays intersect a cell. The number of absorptions a ray experiences depends on whether it arrives before or after previous rays that have intersected a cell. If the timestep is small enough, such that the ionizing front does not move significantly over a timestep, the order rays hit a cell does not matter. Our code is not always operating in this limit. It attempts to reach a converged solution by having many rays from a source intersect a cell that is within ~ 30 Mpc from the QSO and by randomizing the order in which the rays are cast.

2.1. Background

HeII-ionizing photons with $\lambda_{\text{mfp}} > 3.5 L_{\text{box}}$, where λ_{mfp} is calculated assuming a homogeneous IGM with $x_{\text{HeII}} = 1$, are put into a pool of background photons, which we specify with 10 frequency bins.⁶ Five keV is the maximum photon energy that is included in the background. Averaged over a timestep, each cell receives a background flux of

$$f_\nu = c \frac{N_{\text{bk},\nu}}{L_{\text{box}}^3 \tau_\nu} [1 - \exp(-\tau_\nu)], \quad (1)$$

⁶ This number is for the 186 Mpc box. We use the criterion $\lambda_{\text{mfp}} > 1.8 L_{\text{box}}$ for the 429 Mpc box.

where $N_{\text{bk},\nu}$ is the total number of background photons in frequency bin ν , L_{box} is the size of the box, and τ_ν is the average optical depth of a photon traveling a distance $c \Delta t$ in the simulation volume.⁷

2.2. Temperature

Our code tracks the temperature evolution of the gas, which is governed by the differential equation (e.g., Hui & Gnedin 1997)

$$\frac{dT}{dt} = -2 H T + \frac{2 T}{3 \Delta_b} \frac{d\delta_b}{dt} - \frac{T}{\sum \tilde{X}_i} \frac{d \sum \tilde{X}_i}{dt} + \frac{2}{3 k_B n_{\text{tot}}} \frac{dQ}{dt}, \quad (3)$$

where \tilde{X}_i is defined such that the number density in species i is $(1 + \delta) \tilde{X}_i \rho_b / m_p$, dQ/dt is the heating rate, and n_{tot} is the total number of gas particles.⁸ The first and second terms on the right-hand side of equation (3) account for the effect of adiabatic heating/cooling, the third term describes the change in the number of gas particles owing to ionizations and recombinations, and the final term accounts for radiative heating and cooling. The most important process that contributes to dQ/dt is photo-heating, but atomic cooling, recombination cooling, dielectric recombination cooling, collisional cooling,

⁷ A separate background is also included to determine the ionization state of hydrogen. While we could include this ionizing background self-consistently by extrapolating the 4 Ry luminosity of quasars in the box to 1 Ry, there is mounting evidence that there is a substantial stellar contribution to the ionizing background at these redshifts (Madau et al. 1999; Steidel et al. 2001; Boksenberg et al. 2003; Sokasian et al. 2003; Faucher-Giguère et al. 2008b). Therefore, we instead adopt an empirical approach, utilizing observations of the HI photo-ionization rate (Γ_{HI}) from the HI Ly α forest.

Bolton et al. (2005), Becker et al. (2007), and especially Faucher-Giguère et al. (2008b) measured the hydrogen photo-ionization rate (Γ_{HI}) to be flat as a function of redshift using the flux decrement method on $2 < z < 4$ HI Ly α forest data and to within a factor of ~ 2 given by

$$\Gamma_{\text{HI}}(z) = 10^{-12} \text{ s}^{-1}. \quad (2)$$

We use this function for $\Gamma_{\text{HI}}(z)$ throughout this paper and for all simulated redshifts. The background level primarily affects our calculations regarding the HI Ly α forest opacity and does not influence our conclusions.

The approximation of a global background is excellent for tracking the photo-ionization of HI at the redshifts relevant for HeII reionization because the m.f.p. for hydrogen ionizing photons is measured to be $l_{\text{mfp}} = 300 [(1+z)/4]^{-3.2}$ Mpc (Meiksin & White 2004), with present uncertainty at a factor of two level (Faucher-Giguère et al. 2008b). A rough estimate is that there is 1 galaxy per Mpc³ and one quasar per 30³ Mpc³, such that there are thousands of sources that contribute to the background within l_{mfp} . This implies that spatial fluctuations in the background are small at $z \lesssim 5$ (Croft 2004; McDonald et al. 2005). For similar reasons, this approximation is also valid for the hard HeII-ionizing photons that are treated as a global background by our code.

⁸ Several published studies that solve for T use an incorrect form for the third term on the R.H.S. of equation (3) – a factor of $-2/3$ different from what appears in this equation. This mistake stems essentially from writing the first law of thermodynamics as $dq_p = de_p + P d(1/n)$, where q_p is the external heating per particle, e_p is the energy density per particle, and P is the pressure. The correct form for this equation is $dq_p = de_p + P d(1/(\mu m_p n))$, where μ is the mean molecular weight and m_p is the proton mass. Unlike in the former equation, the second term in the correct version does not depend on the ionization state of the gas. This mistake results in the temperature of the gas increasing during ionization processes by too large an amount.

and Compton cooling off of CMB photons are also included in our calculation.⁹ We use an implicit solver to obtain a solution to equation (3).

To evaluate the Lagrangian time derivatives in equation (3) within a cell, the locations of particles at the beginning of the timestep are used to calculate the initial values for the ionization, density, and temperature. Namely, each particle initially has the ionization, density, and temperature of the cell that it was in at the end of the previous timestep. Summing up these previous values for all the particles within the cell (with the appropriate weighting to conserve energy) gives the initial values of these scalars.

The density field is calculated from the gridded N-body particles. The dynamics of N-body particles are different from that of gas particles. Therefore, it is questionable whether this algorithm will be able to predict the correct temperature. However, Hui & Gnedin (1997) demonstrated that the temperature evolution produced from smoothing dark matter simulations at the Jeans scale provides good agreement with the temperature evolution seen in hydrodynamic simulations. In addition, Hui & Gnedin (1997) showed that even using linear perturbation theory to evolve the density field provides reasonable agreement with the evolution of the T - Δ_b relation in hydrodynamical simulations, suggesting that exactly capturing the Jeans scale is not crucial. A comparison of the temperature evolution our code predicts with the Hui & Gnedin (1997) analytic formula is presented in Appendix B.

The simulations in our study capture the Jeans scale to varying degrees. [Although, more appropriate may be the filtering scale, which is typically a factor ~ 2 smaller (Gnedin & Hui 1998).] The Jeans mass is given by

$$M_J = 9.6 \times 10^9 \Delta_b^{-1/2} \left(\frac{T}{10^4 \text{ K}} \right)^{3/2} \left(\frac{1+z}{4} \right)^{-3/2} M_\odot. \quad (4)$$

Compare this mass to the mass within a grid cell in our simulations:

$$M_{\text{cell}} = 1.4 \times 10^{10} \Delta_b \left(\frac{L_{\text{box}}}{186 \text{ Mpc}} \frac{256}{N_{\text{cells}}} \right)^3 M_\odot. \quad (5)$$

These masses are comparable for $\Delta_b = 1$ for our fiducial resolution and box size ($L_{\text{box}} = 186 \text{ Mpc}$ and $N_{\text{cells}} = 256$). For the 512^3 runs in the 186 Mpc box, $M_{\text{cell}} \ll M_J(z=3)$ for $\Delta_b = 1$, and, in the 429 Mpc box, 512^3 resolution achieves $M_{\text{cell}} \approx M_J(z=3)$ for $\Delta_b = 1$. Most important, our conclusions do not change if we use the 512^3 mesh rather than 256^3 .

2.3. Secondary Photo-ionizations and Heating

We assume that the excess energy above E_{HeII} from a photon that photo-ionizes a HeII ion goes into heating the gas. In reality, the secondary electron produced by the photo-ionization may go on to collisionally ionize another atom, excite a bound electron, or suffer coulomb collisions that heat the gas (Shull & van Steenberg 1985).

⁹ We assume a power-law of 1.6 in energy flux to calculate the HI photo-heating. We neglect HeI photo-heating in this calculation. The heat input from HeI photo-heating is $\approx 25\%$ that of the HI photo-heating. Uncertainties in the spectrum of the photo-ionizing radiation result in larger uncertainties in the photo-heating rate than the HeI contribution to it.

However, Shull & van Steenberg (1985) find that 99% of the excess energy of a 3 keV X-ray photon that photo-ionizes a HeII ion goes into heating the gas for a situation in which $\bar{x}_{\text{HI}} = 0.05$, $\bar{x}_{\text{HeI}} = 0.05$, and $\bar{x}_{\text{HeII}} = 0.95$. This percentile will increase for the smaller \bar{x}_{HI} , \bar{x}_{HeI} , and photon energies that are most relevant during HeII reionization. Therefore, the assumption that all the excess energy goes into heating the gas is excellent. Furthermore, Shull & van Steenberg (1985) find that interactions with HeII are unimportant in their calculations. Most of the excess energy is converted into heat because collisions with electrons dominate the energy loss of the secondary particles once hydrogen is reionized.

Finally, Compton cooling of the energetic electrons off of the CMB, which was not included in the calculations of Shull & van Steenberg (1985), is a subdominant energy loss mechanism at $z \sim 3$ for the relevant gas densities and photon energies (Madau et al. 2004).

2.4. Heat Input Estimates

Let us develop an understanding for how much heating is expected from HeII reionization. An HeII ionization front heats up the gas behind it by (e.g., Abel & Haehnelt 1999)

$$k_b \Delta T \approx \frac{2 Y_{\text{He}}}{3(8-5 Y_{\text{He}})} \left(\int_{E_{\text{HeII}}}^{\infty} \frac{dE}{E} \times (E - E_{\text{HeII}}) \sigma_{\text{HeII}}(E) J_i(E) e^{-\tau_E} \right) / \left(\int_{E_{\text{HeII}}}^{\infty} \frac{dE}{E} \sigma_{\text{HeII}}(E) J_i(E) e^{-\tau_E} \right), \quad (6)$$

where J_i is the unabsorbed intensity emitted by the source, $J_i(E_\gamma) \exp(-\tau_E)$ is the incident spectrum, τ_E is the optical depth from the source to the gas parcel, Y_{He} is the primordial helium mass fraction, and σ_{HeII} is the HeII photo-ionization cross section.

For $\tau_E \ll 1$ and using that $\sigma_{\text{HeII}} \sim (E_\gamma - E_{\text{HeII}})^{-3}$, equation (6) becomes

$$\Delta T \approx \frac{2}{3} \frac{E_{\text{HeII}}}{27 k_b (2 + \alpha_{\text{UV}})} = 4500 \left(\frac{3.5}{2 + \alpha_{\text{UV}}} \right) \text{ K}, \quad (7)$$

where α_{UV} is the spectral index of $J_i(E_\gamma)$, $E_{\text{HeII}}/(2 + \alpha_{\text{UV}})$ is the average excess energy of a photon above E_{HeII} , and 27 is the number of particles over which this energy is distributed assuming $Y_{\text{He}} = 0.25$. If τ_E is appreciable, equation (7) still has some relevance, except replace α_{UV} with the effective spectral index of the incident spectrum, which will be harder, resulting in larger ΔT .

The value 4500 K in equation (7) is quite small, and it might be difficult to reconcile such a small ΔT with measurements of ΔT in the Lyman- α forest (assuming HeII reionization is the cause of the temperature increase). It turns out that this number is an underestimate for the total heating during HeII reionization. If we assume all photons up to E_{max} are absorbed then the heat injection is

$$\Delta T \approx 31,000 \left(\frac{0.5}{\alpha_{\text{UV}} - 1} \right) \left(1 - \frac{\alpha_{\text{UV}} E_{\text{HeII}}^{\alpha_{\text{UV}}-1}}{E_{\text{max}}^{\alpha_{\text{UV}}-1}} \right) \text{ K}, \quad (8)$$

where we have kept only the leading order in $E_{\text{HeII}}/E_{\text{max}}$. This equation can be derived by setting $\sigma_{\text{HeII}} = 1$ and $\tau_E = 0$ in equation (6) and performing the integrals. The

heating implied by equation (8) is much larger than by equation (7). Several hundred eV photons have optical depth unity on roughly the 186 Mpc box scale assuming a homogeneous IGM with $x_{\text{HeII}} = 1$, so these photons will be absorbed somewhere in the IGM if they are produced during HeII reionization (a photon travels 850 Mpc between $z = 4$ and 3, roughly the interval over which HeII reionization occurs in our simulations).

If we set $E_{\text{max}} = 350$ eV [which has $\tau \sim 1$ over 200 Mpc for $\bar{x}_{\text{HeII}} = 1$] yields $\Delta T = 13,000$ K using equation (8) and $\alpha_{\text{UV}} = 1.5$ (and 15,000 K if we had not expanded in $E_{\text{HeII}}/E_{\text{max}}$). This number is comparable to the average temperature increase seen in our simulations discussed in Section 4.

2.5. Recombination Rates

The recombination timescale for HeIII is $t_{\text{rec,HeIII}} = [\alpha_{\text{A}}(T) n_e(z)]^{-1}$, where α_{A} is the relevant Case A recombination coefficient and \bar{n}_e is the mean electron density. If $t_{\text{rec,HeIII}}$ is expressed in terms of the age of the Universe, $t_{\text{uni}} \approx 2/3 H(z)^{-1}$, it yields

$$\frac{t_{\text{rec,HeIII}}}{t_{\text{uni}}} \approx 0.6 \left(\frac{T}{10^4 \text{ K}} \right)^{0.7} \left(\frac{1+z}{4} \right)^{-3/2} \Delta_b^{-1}, \quad (9)$$

where Δ_b is one plus δ_b – the overdensity in gas – and we have assumed that the intergalactic hydrogen is fully ionized. The HeII recombination timescale is 5.4 times shorter than this timescale for hydrogen at $T = 2 \times 10^4$ K. This fact is important because additional recombinations require more ionizations to reionize the IGM, which requires more energy injection. Although, the recombination timescale is comparable to the gas cooling timescale, so that the additional heating does not lead to significantly larger temperatures in regions that recombine and are reionized (Bolton et al. 2008b).

Since the number of HeIII recombinations during HeII reionization is large, recombination radiation could contribute significantly to Γ_{HeII} .¹⁰ Unfortunately, HeII-ionizing photons produced by recombinations to the ground state of HeII are not followed by our radiative transfer code. Fardal et al. (1998) showed that this radiation could contribute to the photo-ionization rate at the $\sim 20\%$ level. Tracking this radiation with our code would be prohibitively expensive because it would require treating all cells as sources of ionizing radiation.¹¹ Instead, our code either assumes that this recombination radiation contributes locally to ionizations, which it does by using the Case B recombination coefficient for HeII, or that these photons do not contribute to ionizations of the diffuse IGM, which it does by using the Case A recombination coefficient. The code also uses the corresponding recombination cooling rates for these two

cases. Treating HeII recombinations with the Case A coefficient results in a hotter IGM; more ionizing photons from QSOs are required to ionize the IGM, and the QSO photons are harder than the photons from recombinations to the ground state. Neither Case A nor Case B is correct in detail, and the most appropriate choice for the recombination coefficient depends on the ionization state of the gas.

During HeII reionization, a large fraction of the ionizing photons from recombinations will ionize the diffuse IGM since the HeII bubble size is comparable to the m.f.p. to be absorbed in dense systems. In this case, Case B is the better choice. Since this paper is most interested in studying the HeII reionization process, most of our simulations use the Case B recombination coefficient. However, we have run a case that uses the Case A coefficient (simulation D4 in Table 1). This simulation results in a slightly higher volume-averaged temperature at the end of HeII reionization, with the difference being $\Delta \bar{T} \approx 1000$ K when compared to a similar simulation that uses the Case B rate. The character of the ionization and temperature fields are similar between these two simulations.

2.6. Mean Free Path of Ionizing Photons

At our grid scale, subgrid fluctuations may play an important role in absorbing and hardening the typical spectrum – “filtering” – of the HeII-ionizing radiation. In Appendix A, we investigate how dense clumps filter the radiation, and we quantify how well this effect is captured in our simulations. Systems that have HI column densities $N_{\text{HI}} \sim 10^{15} \text{ cm}^{-2}$ are responsible for filtering HeII ionizing radiation in HeIII regions during HeII reionization. These systems are much less overdense and more diffuse than those that limit the m.f.p. of HI Lyman-limit photons, having overdensities of $\delta_b \sim 10$ at $z \sim 3$ (Schaye 2001), and, therefore, they can be captured with lower resolution simulations than their counterparts for HI Lyman-limit photons, which have $\delta_b \sim 100$. In Appendix D, we show that our gridded density field has some success reproducing the column density distribution of these systems when compared to high-resolution hydrodynamic simulations. However, even though the column density distribution is reproduced, the bias and density of these systems is altered compared to high resolution hydrodynamic simulations. In addition, our radiative transfer code may systematically over-ionize these systems, as described in Appendix D. Therefore, we supplement our calculations with two prescriptions (Methods A and B) to study the effect that dense systems have on filtering the ionizing radiation. These methods are described in detail in Appendix A and briefly here.

Method A assumes that the high-column density HeII systems that are resolved in our simulations are in photo-ionization equilibrium. We argue that this approach is reasonable in Appendix A. Method B uses a Haardt & Madau (1996)-like model that takes the distribution of N_{HI} measured from Lyman- α forest spectra and, given a model for the density of these systems and the local value of the photo-ionization rates in the simulation, infers the HeII column density distribution and, therefore, the opacity in these systems. Further, this model assumes that these absorbers are associated with halos in our simulation that have $m > 7 \times 10^9 M_{\odot}$ and

¹⁰ Ground state recombination radiation results in a line profile of the form $\exp[-(E_{\gamma} - E_{\text{HeII}})/kT]$ for $E_{\gamma} > E_{\text{HeII}}$, such that most photons have excess energies of ~ 1 eV above E_{HeII} for characteristic temperatures of $T \sim 10^4$ K. A photon will travel $80 [(1+z)/4]^{-1/2} [(E_{\gamma} - E_{\text{HeII}})/1 \text{ eV}]$ Mpc prior to redshifting below E_{HeII} . This distance is longer than the m.f.p. for such photons to be absorbed, and, therefore, they will typically be absorbed either in dense clouds or by diffuse intergalactic HeII.

¹¹ Two recent ray tracing codes outlined in Pawlik & Schaye (2008) and Altay et al. (2008) have developed techniques to alleviate this problem.

places the absorbers at the locations of these halos in the simulation box. The reader might be wary of such methods for supplementing the resolution, but we find that these prescriptions have only a minor impact on the final results (Section 4.3).

If much of the intergalactic helium is in HeII, the m.f.p. is often limited by diffuse gas rather than by dense systems. In the limit that diffuse regions near the mean density dominate the opacity, the m.f.p for a photon with energy E_γ is given by

$$\lambda_{\text{mfp}} \approx 5 \bar{x}_{\text{HeII}}^{-1} \left(\frac{E_\gamma}{100 \text{ eV}} \right)^3 \left(\frac{1+z}{4} \right)^{-2} \text{ Mpc.} \quad (10)$$

The value of λ_{mfp} scales strongly with E_γ in this limit. A photon with $E_\gamma = 55 \text{ eV}$ has $\lambda_{\text{mfp}} = 0.8 \text{ Mpc}$ at $z = 3$, whereas one with $E_\gamma = 200 \text{ eV}$ has $\lambda_{\text{mfp}} = 40 \text{ Mpc}$. The average energy of a photon for our fiducial UV power-law index of $\alpha_{\text{UV}} = 1.6$ is 150 eV. However, half of the HeII-ionizing photons have $54.4 < E_\gamma < 84 \text{ eV}$.

3. IONIZING SOURCE MODEL

Quasars are the leading candidates for reionizing HeII. Cooling radiation from massive halos is the next most likely contender (Miniati et al. 2004). However, the large fluctuations in the HeII to HI column density ratios at $z \sim 3$ strongly disfavor cooling radiation as being the dominant source of HeII-ionizing photons. Too many sources of cooling radiation are present within one mean free path to source these large fluctuations (Bolton et al. 2006). Finally, there are other more exotic sources that could potentially ionize HeII at $z \lesssim 6$. Possibilities include an unknown source of $> 4 \text{ Ry}$ radiation that is generating the HeII 1640 Å recombination line observed in the composite spectrum of $z \approx 3$ Lyman break galaxies (Shapley et al. 2003; Furlanetto & Oh 2008)¹², or the bi-products of decaying/annihilating light dark matter particles. These more exotic scenarios are probably also in conflict with observations for the same reason that cooling radiation is disfavored. Our study concentrates on HeII reionization by quasars.

Detailed observations of quasars in the past couple decades have provided many clues into the nature of these extremely luminous objects. We now know that quasars are powered by accretion onto super-massive black holes. Presently, the $z \lesssim 3$ luminosity function of QSOs is well constrained over a few decades in luminosity in both the optical and the X-ray. The clustering of these objects reveals that optically selected quasars reside in halos with $m \approx 10^{12} M_\odot$, independent of their luminosity (Porciani et al. 2004; Croom et al. 2005). Although, clustering is currently measured over only about a decade in luminosity. For the analysis in this study, the biggest uncertainties in modeling QSOs are their lifetimes, τ_{QSO} , the fraction of HeII ionizing photons that escape the environs of a QSO into the IGM as a function of direction, and their spectrum between 4 Ry and 1 keV.

Figure 1 plots the bolometric luminosity function of quasars for four redshifts. The black, solid curves are the




FIG. 1.— Comparison of models and measurements of the bolometric QSO luminosity function at 4 redshifts. The data points are from a compilation of observations across different wavebands (Hopkins et al. 2007), using a spectral template to infer the bolometric luminosity function. The solid curves are the Hopkins et al. (2007) best-fit luminosity function (which the quasars in Method I are drawn from) and the dashed curves are the luminosity functions generated with Method II. The solid horizontal lines represent the number densities at which 1 QSO is in the 186 and 429 Mpc boxes.

Hopkins et al. (2007) best-fit model.¹³ The cyan points with error-bars are infrared observations, the green are optical, the blue are soft X-ray, and the red are hard X-ray. The references for these observations are given in Hopkins et al. (2007), and the symbols follow the same conventions as described in their Table 1. Observations in different bands are included in this figure by assuming an intrinsic spectral template and an average column density distribution for the obscuring matter that is motivated by X-ray measurements of N_{HI} . At $z > 3$, the luminosity function falls in amplitude. At $z > 5$ there is not much data to constrain the models. The abundance of QSOs between $3 \lesssim z \lesssim 4.5$ shapes the duration and morphology of reionization in our simulations. Insufficient HeII-ionizing photons are produced at higher redshifts to substantially ionize the HeII with the best-fit luminosity function model. In this model, ~ 2 HeII ionizing photons per helium atom are produced by $z = 3$, about enough to reionize the HeII.

The faint-end slope of the the luminosity function $d\phi/d\log L$ is constrained to be rather flat at $z \approx 3$, with power-law index -0.5 . This flatness means that most of the HeII-ionizing photons are produced by $L \sim L_*$ quasars. The horizontal lines in the top panel in Figure 1 demark the column densities that yield 1 quasar in the 186 and 429 Mpc simulation boxes. Our simulations are large enough to contain many L_* QSOs.

Quasars are placed in our simulations with one of two methods:

3.1. QSO Method I

The lifetime of all quasars is fixed to be τ_{QSO} , and quasars are placed in halos of mass $m \propto L^{4/3}$. This

¹² Furlanetto & Oh (2008) find that only if HeII ionizing photons can escape with $f_{\text{esc,HeII}} \gtrsim 0.5$ from galaxies can they contribute a significant fraction of the ionizing background. These values are much larger than the theoretical expectation for $f_{\text{esc,HeII}}$ from galaxies (e.g., Gnedin et al. 2008).

¹³ The Hopkins et al. (2007) luminosity function is derived by simultaneously fitting to the infrared, optical, and X-ray luminosity functions. These fits model the observed quasar spectra with a bolometric luminosity-dependent intrinsic spectral template plus a model for the column-density distributions in N_{HI} and dust. This procedure results in a good fit to the observed luminosity function in all bands (Hopkins et al. 2007).

scaling assumes QSOs shine at their Eddington Luminosity ($L_{\text{edd}} \propto M_{\text{BH}}$) with some duty cycle and that $M_{\text{bh}} \sim \sigma^4 \sim m^{4/3}$.¹⁴ The number of QSOs in the box at a given luminosity is drawn from a Poisson distribution with mean set by the observed luminosity function. We use the Hopkins et al. (2007) luminosity function evaluated at 912 Å and extrapolate across the relevant band with a power-law α_{UV} .

3.2. QSO Method II

Recently, simulations of merging galaxies have been used to predict quasar light curves (Hopkins et al. 2005, 2006), and these predictions have been remarkably successful at reconciling observations of the quasar luminosity function, quasar clustering, and the unresolved X-ray background (e.g., Hopkins et al. 2006). The bolometric light curve of the quasars in merger simulations can be roughly parametrized as

$$L_{\text{bol}}(t) = 0.8 L_{\text{edd}}(m_{\text{bh}}) \times \begin{cases} \exp[t/\tau_S] & t < 0 \\ (1 + t/\tau_S)^{-b} & t > 0, \end{cases} \quad (11)$$

where m_{bh} is the super-massive black hole mass, $b = 1.7 + 0.7 \times \log_{10}(L_{\text{edd}}/10^{12} L_{\odot})$, L_{edd} is the Eddington luminosity, and $\tau_S = 40$ Myr. To map from $L_{\text{bol}}(t)$ to the intrinsic luminosity at the HI Lyman-limit, we take the Hopkins et al. (2007) L_{bol} -dependent spectral template.

Next, equation (11) plus a relationship between halo mass and $L_{\text{edd}}(m_{\text{bh}})$ allows one to deconvolve from the observed QSO luminosity function the triggering rate per unit volume of quasars as a function of time. The triggering rate is defined as the rate quasars shine at their peak luminosity. We use the Magorrian relationship to relate a galaxy's stellar mass to its super-massive black hole mass (Hopkins et al. 2007a,b), and the halo occupation model of Tinker et al. (2005); Conroy et al. (2006); Vale & Ostriker (2006) (for details, see Hopkins et al. (2008a)) to map from halo mass to stellar mass.¹⁵ Figure 1 show the luminosity function that this model generates (red, dashed curves). While it differs somewhat from the best-fit luminosity function (black, solid curves), it agrees well with the observations. To implement this model in our simulations, we switch on quasars in the simulation box using the triggering rate per unit halo mass described above, accounting for Poisson scatter.

3.3. Discussion of QSO Models

Figure 2 shows the HeIII bubble size probability distribution function (PDF) in Method I and Method II. These bubble sizes correspond to the size of a bubble if all the QSOs sit in separate spherical HeIII regions and there are no recombinations. All curves are calculated assuming $\alpha_{\text{UV}} = 1.6$. For the thin curves, the PDF is volume-weighted, and, for the thick curves in the bottom panel,

¹⁴ The specific mapping we use is $m = 7.9 \times 10^{12} [(5\nu L_{\nu})/10^{46} \text{ erg s}^{-1}]^{3/4} [(1+z)/4]^{1.5} M_{\odot}$, where νL_{ν} is evaluated at 1 Ry. Our results depend weakly on the mapping between L and m because the morphology of HeII reionization turns out to be determined primarily by Poisson fluctuations in the abundance of quasars rather than by their clustering.

¹⁵ As a final ingredient, since there is scatter in the relationship between \bar{m}_{bh} and M_{stellar} , we assume that the actual m_{bh} is log-normal function with dispersion 0.3 dex and logarithmic mean given by the Magorrian relationship. This step is necessary in order to reproduce the bright end of the luminosity function.




FIG. 2.— HeIII bubble size distributions in our two QSO models assuming $\alpha_{\text{UV}} = 1.6$, isotropic emission, 1 QSO per bubble, and no recombinations. The black solid curves correspond to $z = 2$, the green dashed to $z = 4$, and the blue dotted to $z = 6$. Top panel: Thin curves are the volume-weighted probability distribution of bubble radii in Method I. These curves assume $\tau_{\text{QSO}} = 40$ Myr. Bottom panel: Thick curves are the volume-weighted probability distribution in Method II, and the thin curves are the number-weighted probability distribution.

it is number-weighted. Method I has an infinite number of quasars with $L < L_*$ and so the number-weighted PDF is not well-defined.

The HeIII bubbles from the quasars in Method II have a characteristic size of 35 Mpc in the volume-weighted PDF, and the probability distribution of R_b does not vary significantly with redshift. With Method I, the characteristic radius is a bit smaller, 20 Mpc for $\tau_{\text{QSO}} = 40$ Myr, and the dispersion in R_b is larger. Method I has a non-negligible probability for extremely large bubble sizes to exist, bubbles with $R_b \gtrsim 70$ Mpc. These large bubbles are averted in Method II because brighter quasars have the shorter lifetimes in this method. The properties of HeII reionization do not change substantially if we employ Method I or Method II in our simulations.

Given the large observational uncertainty in the mean and distribution of α_{UV} (see Appendix E), we leave these quantities as free parameters in both QSO Models. For each quasar in the simulation box, α_{UV} is drawn from a Gaussian distribution with mean $\bar{\alpha}_{\text{UV}}$ and s.d. σ_{α} . Our fiducial model assumes $\bar{\alpha}_{\text{UV}} = 1.6$ and $\sigma_{\alpha} = 0.2$, which is motivated by the results of Telfer et al. (2002) and Steffen et al. (2006) [Appendix E]. We find that the amount of heating from HeII reionization is sensitive to the spectrum of QSOs.

We also include a parameter ζ , a suppression factor for the number of HeII-ionizing photons that are emitted by QSOs. This parameter gives us a knob to tune in order to have HeII reionization occur at a desired redshift, and, given current uncertainties in the QSO luminosity function, can be adjusted at a factor of 2 level.

4. SIMULATIONS

We have run a series of simulations to investigate the morphology of HeII reionization. All of our simulations track quasars to $L_{\text{bol}} > 10^{43}$ or $10^{44} \text{ erg s}^{-1}$, approximately 3 decades below L_* . Table 1 lists the simulations

TABLE 1
HeII REIONIZATION SIMULATIONS.

Sim. ^a	L_{box} (Mpc)	LLS ^b	QSO Model	Emission ^c	ζ	Grid	Comments
D1	186		II	iso	0.75	256	
D2	186		II	iso	0.75	256	no density fluctuations ($\delta_b = 0$)
D3	186		II	iso	0.75	256	twice as many frequency bins as D1
D4	186		II	iso	0.75	256	case A recombinations for HeII
L1	186	A	II	iso	0.75	256	
L1b	186	A	II	iso	0.75	256	$\gamma = 0$ at $z = 6$
L1c	186	A	II	iso	1.5	256	slightly earlier reionization than L1
L1d	186	A	II	iso	0.75	256	extremely conservative photon termination
L2	186	A	II	iso	1	512	high resolution run
L3	186	B	II	iso	0.75	256	includes $N_{\text{HI}} > 10^{14.5} \text{ cm}^{-2}$
X1	186	A	II	iso	0.7	256	omits photons in background
S1	186	A	II	beamed	0.75	256	beam uses Gilli et al. (2007) model
S2	186	A	I	iso	2	256	$\tau_{\text{QSO}} = 10 \text{ Myr}$
S3	186	A	II	iso	0.75	256	assumes $\bar{\alpha}_{\text{UV}} = 1.2$
S4	186	A	II	iso	0.75	256	assumes $\bar{\alpha}_{\text{UV}} = 0.6$
S4b	186	A	II	iso	1.25	256	assumes $\bar{\alpha}_{\text{UV}} = 0.6$, no background
B1	429		II	iso	1	256	
B2	429	A	II	iso	1	512	

^a Unless specified otherwise, the initial conditions for these simulations are $\gamma - 1 = 0.3$ and $T_0 = 10^4 \text{ K}$ at $z = 6$, and $\bar{\alpha}_{\text{UV}} = 1.6$. All simulations allow for a random dispersion in α_{UV} with $\sigma_\alpha = 0.2$. At fixed ζ , the number of photons per timestep is the same in all the simulations that use QSO Model II, independent of $\bar{\alpha}_{\text{UV}}$.

^b The “LLS” column specifies the method for capturing the dense absorbers, as discussed in Appendix A.

^c iso = isotropic emission




FIG. 3.— Volume-averaged global history of Γ_{HeII} in s^{-1} , T in kilo-Kelvin, and x_{HeIII} , in six of our simulations.

discussed in this paper, and Figure 3 plots the volume-averaged history of Γ_{HeII} , T , and x_{HeIII} in six of these simulations. The simulations in Figure 3 use different prescriptions to include the QSOs and to capture the high column-density absorption systems. Yet, with a minor amount of tuning, the end of HeII reionization occurs

near $z \sim 3$ in all of the simulations.¹⁶ The known population of quasars is able to ionize the HeII in the Universe by $2.5 < z < 3.5$, approximately the redshift range where several observations suggest that HeII reionization is ending. In the end, 2.2 photons per helium atom are required to reionize the HeII to $x_{\text{HeII}} = 0.95$ in simulation D2 (which employed no filtering method) and 3.4 photons per helium atom are required in simulations L1 and L3 (which used filtering methods A and B).

The middle and top panels in Figure 3 show the history of \bar{T} and $\bar{\Gamma}_{\text{HeII}}$. In these simulations the temperature of the IGM peaks near the end of HeII reionization, and $\bar{\Gamma}_{\text{HeII}}$ increases during this process. The jaggedness of $\Gamma_{\text{HeII}}(z)$ owes to sample variance.¹⁷

The remainder of this section discusses our simulations in detail. But first, it is important to highlight a drawback of our radiative transfer method. Time steps are set to $\Delta t = 10 \text{ Myr}$, such that a ray will travel a distance $c\Delta t$ over a timestep before being held in memory. This introduces a characteristic scale of $c\Delta t = 3$ proper Mpc that can appear in the simulations. In particular, the effect of a finite timestep is sometimes apparent in the Γ_{HeII} field, through rings that show sharp gradients in Γ_{HeII} . These small artifacts do not affect our conclusions.

¹⁶ Note that we tune the parameter ζ , which normalizes the total ionizing emissivity, between 0.75 and 2 in the simulations presented here in order for HeII reionization to end at $z \approx 3$. Given observed uncertainty in the QSO luminosity function and spectral behavior near the HeII Lyman-limit, such tuning is acceptable. This aids comparison between simulations, and it also allows the simulations to address better the set of observations that indicate that HeII reionization is nearing completion at $z \approx 3$.

¹⁷ The jaggedness is particularly acute in simulation S2 because it uses the shortest quasar lifetimes, it has the fewest quasars at high redshifts, and its brightest quasars have the same lifetime as its dimmest ones.

4.1. Homogeneous Universe

Simulation D2 is run in the 186 Mpc box using the positions of the halos in the N-body code and Method II for populating these halos with quasars. In order to isolate some interesting physical effects, this simulation is run in a homogeneous density field ($\Delta_b(\mathbf{x}, z) = 1$) rather than using the density field from the N-body simulation as in the other cases. The left-most panels in Figure 4 depict the HeII ionization field at several different times. Initially, the HeIII bubbles around each QSO are isolated spheres with characteristic circular radii of ≈ 20 Mpc.¹⁸

The edges of the HeIII regions are not sharp, with the region over which $0.1 < x_{\text{HeIII}} < 0.9$ spanning several Mpc and with a lower ionization tail extending much further. When the Universe is roughly a few tenths ionized, the spherical HeIII regions begin to merge and form larger non-spherical HeIII regions. In addition, the fractionally ionized regions that extend far ahead of the ionizing front overlap, creating a global ionization floor. By the time $\bar{x}_{\text{HeIII}} > 0.8$ in simulation D1, $x_{\text{HeIII}} > 0.2$ everywhere.

Since the recombination timescale at mean density is shorter than the Hubble time and comparable to the HeII reionization timescale, relic HeIII regions recombine significantly over the course of HeII reionization. Relic HeIII regions in which the HeIII has substantially recombined are present in the top-right panel in Figure 4. However, these regions persist only while the HeII ionizing background has $\Gamma_{\text{HeII}} \lesssim \alpha_B \bar{n}_e \approx 2 \times 10^{-17} \text{ s}^{-1}$ at $z \approx 4$, where α_B is the Case B recombination coefficient. Otherwise, ionizations will be sufficient to balance recombinations. Relic HeIII regions that have substantially recombined exist in our simulations primarily during the early phases of HeII reionization ($\bar{x}_{\text{HeIII}} \lesssim 0.5$). The fourth column in Figure 4 displays Γ_{HeII} . The condition $\Gamma_{\text{HeII}} > 2 \times 10^{-17} \text{ s}^{-1}$ is met in most regions in the IGM at times when $\bar{x}_{\text{HeII}} \gtrsim 0.5$.

The panels in the second column of Figure 4 show \log_{10} of the transmission in the HeII Ly α forest. For a homogeneous Universe, there is little transmission in the HeII Ly α forest, even at times near the end of HeII reionization. The inclusion of density inhomogeneities significantly enhances the transmission level (Section 4.2).

The IGM temperature field has a different morphology from the x_{HeIII} field (compare the third column in Fig. 4 with the first). Early on in HeII reionization, the morphologies of both the x_{HeIII} and T fields are driven by isolated HeIII regions. The temperature of an isolated HeIII region is $\Delta T \approx 7,000$ K above the average temperature in HeII regions, with hotter regions at the HeIII front edge and slightly colder regions inside the front. Once ionized, the HeIII region begins to cool adiabatically (owing to cosmic expansion). Regions that were ionized earlier are cooler than regions that have been ionized more recently. The volume-averaged temperature increase of 10,000 K during HeII reionization is

¹⁸ This value for R_b may seem somewhat smaller than those projected in Figure 2 for several reasons: (1) A random slice through a bubble will on average have a factor of 2 smaller radius than the 3-D radius; (2) active quasar bubbles are smaller than the final bubble size – the quantity plotted in Figure 2 (the quasars in Method II have long lifetimes that are comparable to the Salpeter time); (3) a fraction of the ionizing photons are deposited ahead of the front and some reionize recombining atoms; (4) $\zeta < 1$ in D2.

consistent with estimates that assume that the IGM absorbs all photons with energies less than a few hundred eV during HeII reionization (Section 2.3).

The right-most column in Figure 4 displays the cumulative amount of HeII photo-ionization heating measured in Kelvin. In the temperature panels it is apparent that the coolest regions are the first regions to be ionized, which we justified by the fact that these regions had more time to cool. However, in these photo-heating panels, we see that the regions that are ionized first also have the least heat injection from HeII photo-heating. Regions that are ionized at later times have absorbed more of the hard radiation background before being fully ionized.

4.2. Density Fluctuations

Figure 5 is the same as Figure 4, but features simulation D1 which includes density inhomogeneities using the 256^3 gridded N-body field. The small-scale structure of HeII reionization is changed significantly by including density inhomogeneities (compare the left-most panels in Figures 4 and 5). However, the large-scale morphology of HeII reionization is similar between the homogeneous and inhomogeneous cases.

The structure of the temperature fluctuations is affected the most by density inhomogeneities compared to the other quantities shown in Figure 5. Adiabatic heating and cooling from structure formation imprints additional small-scale features into the temperature field in simulation, and it makes more regions cooler in D1 than in D2 because most of the volume in the IGM is underdense and expanding faster than if it were in the Hubble flow.

The second column in Figure 5 shows the total HeII Ly α forest transmission. The transmission level is significantly higher when density inhomogeneities are included. Underdense regions are responsible for most of the transmission in the $z \approx 3$ HeII Ly α forest. Section 5.3 shows that simulation D2 is consistent with measurements of the HeII Ly α forest mean transmission.

Figure 6 presents simulation B1, which has the same run specifications as simulation D1, but in the larger 429 Mpc box. The larger box provides a better sample of the structures during HeII reionization. On these larger scales, QSO clustering is evident in the maps.

4.3. The Effect of Dense Clumps

Figure 7 compares three simulations that employ different methods for including the effect of dense clumps on filtering the ionizing radiation. Each row features a snapshot from its respective simulation with $\bar{x}_{\text{HeIII}} = 0.85$. The top panels are from simulation D1 (which does not employ an additional filtering method), the middle panels are from simulation L1 (filtering Method A), and the bottom panels are from simulation L3 (Method B). All three simulations are run on the 256^3 grid using quasar Model II. Simulation D1 reaches $\bar{x}_{\text{HeIII}} = 0.85$ at $z = 3.3$, and simulations L1 and L3 approach this mark at slightly later times, at $z \approx 3.15$, owing to the additional filtering. For the most part, the x_{HeIII} and T fields show similar patterns in the three simulations, which suggests that our results depend weakly on the filtering method. The volume-averaged temperatures in the three simulations are within 2000 K of each other at fixed \bar{x}_{HeII} . Furthermore, the distribution of temperatures at mean density




FIG. 4.— Slices through simulation D2, which assumes a homogeneous IGM ($\Delta_b = 1$). Each slice has breadth 186 Mpc and width 3.5 Mpc (except for the second column, which has a width of 1 grid cell or 0.7 Mpc). From top to bottom, the panels in each row feature snapshots at $z = 4.3, 3.5, 3.2$ and 3 respectively, corresponding to $\bar{x}_{\text{HeIII}} = 0.1, 0.5, 0.8$ and 0.99 . The first column shows the spatial distribution of x_{HeIII} . The second column depicts \log_{10} of the transmission in the HeII Ly α forest. The middle column shows the temperature of the IGM. The fourth column plots \log_{10} of Γ_{HeII} , and the right-most column shows the cumulative amount of heat deposited by HeII photo-ionizations. White regions represent higher values than shown in the scale and black represent lower values.

are similar between these simulations, but with the distribution in L1 extending to a few thousand Kelvin higher temperatures than those in the other simulations.

The right-most column in Figure 7 illustrates how the filtering methods affect the m.f.p. of ionizing photons, where the m.f.p. is tabulated by averaging the distance that absorbed photon travel (where this average is weighted by the energy it deposits). The average m.f.p. decreases slightly from simulation D1 to L1 as expected, and again from L1 to L3.

It is surprising that filtering from dense systems does not have a more significant impact, particularly on the temperature fields, which are sensitive to the hardness of the ionizing radiation. If the bubble sizes were much smaller than the photon m.f.p., then it would not be surprising that filtering by dense regions is unimportant – radiation would be filtered minimally by dense clumps prior to reaching diffuse neutral gas. However, the bubble sizes are comparable to the m.f.p. (Appendix A.), and we know that about 1/3 of the photons are absorbed in dense systems in L1 and L3 since HeII reionization in these simulations requires 3.4 photons per HeII ion as

opposed to 2.2 in simulation D1. If we assume that the filtering removes the bottom third of the HeII-ionizing photons in energy and that the IGM is ionized by the remaining, harder radiation, the temperature increase of the IGM would be a factor of 2.4 larger for $\bar{\alpha}_{\text{UV}} = 1.6$ than if the filtering had not occurred (at least if ΔT is computed in the optically thin limit; eqn. 6).

Why is filtering in our simulations different from this toy scenario? The answer is that not just the lowest energy photons are filtered. The effective optical depth of dense systems scales as $\sim E_{\gamma}^{-1.5}$ if the N_{HeII} distribution scales as -1.5 (Zuo & Phinney 1993). This absorption contrasts with the absorption from diffuse neutral gas, for which the effective optical depth will scale closer to $\sim E_{\gamma}^{-3}$. Therefore, the spectral shape of the ionizing radiation is not as altered by filtering by dense systems as by diffuse IGM gas, and it is certainly not as hardened by dense systems as in the above toy model. Furthermore, equation (7) shows that the amount of heating depends weakly on the spectral index of the incident spectrum.

Much of the filtering that hardens the radiation field in the HeII photo-heating panel in Figure 7 occurs in




FIG. 5.— Same as Fig. 4 but using simulation D1 (which includes density fluctuations). The panels in each row are chosen to match the \bar{x}_{HeIII} fractions quoted in the caption of Fig. 4.

the diffuse gas rather than in dense clumps. This is why the temperatures in simulation D2, for which there are no density fluctuations, are similar to the other simulations. Even in the HeIII bubbles in Figure 7, there are significant amounts of neutral gas which do not reside in very overdense locations. Furthermore, photons that penetrate into neutral regions will typically encounter a higher total HeII column density in diffuse gas than in dense systems. These photons are responsible for heating the hottest regions (which are the last regions to be ionized).

Figure 8 compares the globally averaged spectrum in simulations D1 (no filtering) and L1 (filtering Method A). This global average is computed by tabulating the incident spectrum in randomly selected cells in the simulation volume and then averaging. (The global spectrum in L3, which is not shown, is similar to this spectrum in L1.) Reionization ends ($\bar{x}_{\text{HeIII}} = 0.95$) at $z = 3.3$ in D1 and at $z = 3.1$ in L1. The average spectrum has a similar shape, particularly at the beginning of HeII reionization when the impact of filtering is minimal. At later times, the difference is a factor of $\lesssim 2$, and the slope is not significantly changed. Near the end of HeII reionization, the amplitude of the spectrum evolves by order unity in both simulation D1 and L1 on timescales of $\Delta z \approx 0.2$.

This fast evolution suggests that diffuse gas is playing a role in limiting the mean free path in these simulations even at the end of HeII reionization.

After HeII reionization, the filtering method affects the spectrum more because the m.f.p. is limited by dense systems and not by diffuse gas. The value of $\bar{\Gamma}_{\text{HeII}}$, which in a large part determines the transmission in the HeII Ly α forest, is proportional to λ_{mfp} . The photo-ionization rate at the end of the simulation is $2 - 4 \times 10^{-15} \text{ s}^{-1}$ in L1 and L2 whereas it is $7 \times 10^{-15} \text{ s}^{-1}$ and increasing steadily when simulation D1 was terminated at $z = 2.9$.

4.4. Soft X-ray Photons

Simulation X1 is the same as L1 except that photons with m.f.p. larger than 3.5 times box size, background photons with roughly $E_\gamma \gtrsim 500 \text{ eV}$, are not included in the calculation. How much heating and ionizations do these high-energy photons contribute? When the Universe is 90% ionized in HeII, model X1 is 1400 K cooler than model L1. Furthermore, the fraction of ionizations that arise from photons with $E_\gamma \gtrsim 500 \text{ eV}$ is small. The heating from the background is largest in the regions that are ionized last, the regions that are exposed to the background for the longest period of time. The background heating adds an additional several thousand Kelvin in these regions.




FIG. 6.— Slices through simulation B1, which employs the same run specifications as simulation D1 (Fig. 5), but uses the 429 Mpc box. Each panel shows a $429 \times 429 \times 8.4$ Mpc slice through this simulation volume. From top to bottom, $\bar{x}_{\text{HeIII}} = 0.1, 0.5$, and 0.9 .

The relative unimportance of heating from photons with $E_\gamma \gtrsim 500$ eV is reassuring, because we do not model the obscured quasar contribution to the ionizing background. Models of the obscured contribution to the diffuse X-ray background find that obscured quasars (defined as having QSOs with obscuring columns of $N_{\text{HI}} > 10^{21} \text{ cm}^{-2}$) begin to dominate the background above approximately 2 keV (e.g., Gilli et al. 2007).

4.5. QSO Model

The major uncertainties associated with modeling QSOs are their lifetimes, spectra, and whether their emission is beamed. Figure 9 considers 3 QSO models, to illustrate how these uncertainties influence our results. From top to bottom, the panels represent snapshots with $\bar{x}_{\text{HeIII}} \approx 0.1, 0.75$, and 0.9 . Beamed emission (see simulation S1 in Fig. 9 which uses the beaming model described in Appendix E) enhances the structure of the ionization field earlier on in HeII reionization. Similarly a light bulb model with $\tau_{\text{QSO}} = 10$ Myr (simulation S2 in Fig. 9) results in more relic HeIII regions compared to our fiducial case and in slightly smaller bubbles. This fact owes to the shorter quasar lifetime compared to the fiducial case.

The IGM temperatures in S4 and S4b, which have a harder UV spectral index of 0.6 compared to the fiducial value of 1.6 (chosen to better agree with the results of Scott et al. 2004), are significantly higher than the other simulations analyzed in this paper. This result agrees with that of Tittley & Meiksin (2007), who also find a strong dependence of the temperature on the spectral index of the sources. In S4 (right panels in Fig. 9), the hottest regions are typically the least ionized. In the other simulations featured in Figure 9, the hottest regions are typically the most ionized until $\bar{x}_{\text{HeIII}} \gtrsim 0.8$.

Why is the morphology of the temperature so different in S4 whereas the \bar{x}_{HeIII} field is similar to the other simulations (except for a higher ionization floor)? For $\bar{\alpha}_{\text{UV}} = 0.6$, the average energy photon that is absorbed is much higher than in the fiducial model, injecting more heat. However, near quasars it is still the softest photons that are absorbed, resulting in similar structures in the x_{HeIII} field compared to the fiducial model. The hardest photons (which are more plentiful in S4) are absorbed in the neutral regions outside of the HeIII bubbles, leading to the inverted behavior between x_{HeIII} and T .

4.6. Power Spectrum

We have seen that 50 Mpc fluctuations in the temperature and x_{HeIII} are present during and after HeII reionization. These fluctuations modulate the level of absorption as well as the amount of small-scale power in the HI Ly α forest. Here, we use power spectrum statistics to quantify the properties of these spatial fluctuations.

The curves in Figure 10 are calculated from simulation B1, which uses the 430 Mpc box, and represent the power spectrum of x_{HeIII} (top panel), $\delta T \equiv T/\bar{T} - 1$ (middle panel), and the cross correlation coefficient of x_{HeIII} and δ_b (bottom panel) from simulation D1, which is defined as $r \equiv P_{\delta_{x_{\text{HeIII}}}\delta_b} / [P_{\delta_{x_{\text{HeIII}}}\delta_{x_{\text{HeIII}}}} P_{\delta_b\delta_b}]^{-1/2}$, where $\delta_{x_{\text{HeIII}}} \equiv x_{\text{HeIII}}/\bar{x}_{\text{HeIII}} - 1$. We have checked that the power spectrum in B1 agrees well with that in D1, which is identical to B1 except run in the 429 Mpc box. The only scales where the two simulations do not agree correspond to $k > 2 h \text{ Mpc}^{-1}$, where there is an upturn owing to shot noise in B1 that is not present on these scales in D1.

The x_{HeIII} fluctuations peak on ~ 50 Mpc scales, even early in the HeII reionization process when $\bar{x}_{\text{HeII}} \approx 0.1$, with the peak in the amplitude of the fluctuations occurring at $\bar{x}_{\text{HeIII}} \approx 0.5$. This peak scale is remarkably constant throughout HeII reionization. This contrasts with the evolution of the power spectrum in most models of hydrogen reionization, in which the growth of large HII regions around clusters of sources leads to the power shifting to larger scales with increasing \bar{x}_{HI} (e.g., McQuinn et al. 2007).

The fluctuations in the temperature field are shown in the middle panel in Figure 10 at several \bar{x}_{HeIII} . For comparison, the thin blue dotted curve is from a simulation where HeII reionization does not occur at $z < 6$ and that is initialized to have $\gamma - 1 \approx 0$ and $T = 10^4$ K at $z = 6$. This curve is at the same redshift, $z = 4.0$, as the $\bar{x}_{\text{HeIII}} = 0.5$ curve. This illustrates that HeII reionization produces temperature fluctuations on much larger scales than structure formation.

At $k \lesssim 0.3 \text{ Mpc}^{-1}$, the fluctuations from patchy HeII




FIG. 7.— Three rows use snapshots at $\bar{x}_{\text{HeIII}} = 0.85$. The top panels are from a simulation that does not employ an additional filtering method, the middle panels are from a simulation that uses filtering method A, and the bottom panels are from a simulation that uses filtering Method B.




FIG. 8.— Globally averaged radiation spectrum at various times during our simulations, computed by tabulating the incident spectrum in many randomly selected cells and then averaging. The thin curves are for simulation D1 and the thick curves are for L1. L1 differs from D1 in that it uses filtering method A.

reionization are the dominant source of the T fluctuations. The scale of the peak in these fluctuations is comparable to the HeII bubble sizes early on. However, when $\bar{x}_{\text{HeIII}} = 0.3, 0.5$ and 0.97 , large fluctuations are present on all scales captured in our box. Interestingly, the amplitude of large-scale temperature fluctuations are largest at the beginning of HeII reionization, with $(\Delta T/\bar{T})_{\text{rms}} \equiv [k^3 \Delta T / (2\pi^2 T)]^{1/2} = 0.2$. Fluctuations in temperature are observable in the HI Ly α forest (Zaldarriaga 2002; Theuns et al. 2002a).

The degree of correlation between the ionization and density fields reveals how important QSO clustering is in shaping the morphology of HeII reionization. The

bubbles begin to correlate with the density field on 50 Mpc scales (bottom panel in Fig. 10). On smaller scales, x_{HeIII} is anti-correlated with δ_b . The level of anti-correlation increases with \bar{x}_{HeIII} as the HeIII regions encompass more dense neutral systems.

5. OBSERVATIONAL IMPLICATIONS

This section discusses how HeII reionization affects observations of the IGM. The observables that are addressed are the IGM temperature, the T - Δ_b relation, and the mean transmission in the HI and HeII Ly α forests, all of which can be estimated from high-redshift quasar spectra.

5.1. Temperature Evolution

Figure 11 plots the evolution of the temperature PDF in simulations D2, L1, and S3, tabulated from grid cells with $-0.1 < \delta_b < 0.1$. Simulation D2 has no density fluctuations and $\bar{\alpha}_{\text{UV}} = 1.6$; simulation L1 includes density fluctuations, $\bar{\alpha}_{\text{UV}} = 1.6$, and filtering Method A; and simulation S3 is the same as L1 but with $\bar{\alpha}_{\text{UV}} = 1.2$.

The evolution of the PDF in simulation L1 is characteristic of our simulations with $\bar{\alpha}_{\text{UV}} = 1.6$. At times in L1 for which $\bar{x}_{\text{HeIII}} \ll 1$, most of the IGM has $T \approx 1 \times 10^4$ K, except ionized regions which have $T \approx 1.6 \times 10^4$ K. The result is a bimodal temperature PDF (see the $\bar{x}_{\text{HeIII}} = 0.1$ curve in the middle panel in Fig. 11). As HeII reionization proceeds, the PDF shifts to higher temperatures. If only the insides of HeIII regions were heated, this PDF would become broader rather than move towards higher T in the manner seen in Figure 11, implying that hard photons are streaming far from their sources prior to being absorbed. The bimodality of the PDF decreases with increasing \bar{x}_{HeIII} until




FIG. 9.— Morphology of HeII reionization for three possible source models. In S1 the emission from the quasars is beamed as described in the text, in simulation S2 the quasars have a light bulb behavior with $\tau_{\text{QSO}} = 10$ Myr, and in simulation S4 the quasars have a hard spectrum with $\bar{\alpha}_{\text{UV}} = 0.6$. From top to bottom, the panels represent snapshots with $\bar{x}_{\text{HeIII}} \approx 0.1, 0.75$, and 0.9 . *Note that temperature panels for simulation S4 have an adjusted range that extends to 40,000 K compared to 25,000 K for the other simulations.*

$\bar{x}_{\text{HeIII}} = 0.6$, at which point the PDF is fairly Gaussian. For $\bar{x}_{\text{HeIII}} > 0.6$, the PDF again becomes skewed as the harder photon background builds up and starts to appreciably heat regions far from quasars. The temperature PDF at the end of HeII reionization is broad, extending between 12,000 K and 35,000 K in simulation L1. The temperature of a gas parcel at the end of HeII reionization is a result of many factors, including the redshift(s) it was ionized and the hardness of the radiation field that ionized the parcel.

The evolution of the temperature PDF in simulations D2 and S3 are similar to that in L1. The PDF in D2 does not extend to as high temperatures as in L1. This implies that density inhomogeneities contribute to the temperature of the hottest gas parcels. The evolution of the PDF is also similar in S3 to L1, except the harder spectrum in S3 results in higher temperatures.

Schaye et al. (2000) and Ricotti et al. (2000) estimated the temperature of the IGM from the widths of the narrowest lines in the HI Ly α forest, calibrating their measurements with numerical simulations that essentially assumed a power-law $T\text{-}\Delta_b$ relation at relevant Δ_b . These studies claimed to have detected a sudden increase in the IGM temperature between $z \approx 3.5$ and 3. A subsequent analysis using a similar methodology by McDonald et al. (2001) did not detect a sudden increase in temperature, but rather a temperature at mean density of $T_0 \approx 17,000 \pm 2000$ K at $z = 2, 3$, and 4. Zaldarriaga et al. (2001) derived similar num-

bers to McDonald et al. (2001) from measurements of the HI Ly α forest power spectrum. Although, both McDonald et al. (2001) and Zaldarriaga et al. (2001) measured the temperature in three coarse bins centered at $z = 2.4, 3.0$ and 3.9 , possibly obscuring a sudden temperature change.

The bottom panel of Figure 12 plots the Schaye et al. (2000), Ricotti et al. (2000), and McDonald et al. (2001) measurement values for T_0 as well as the evolution of this quantity in simulations L1, D1, and S4b (thick curves). In general, our calculations over-predict the measured values of T_0 . However, since these observations look for the narrowest lines in the forest, it is probable that they are most sensitive to the coolest temperatures at a given density as is argued in Furlanetto & Oh (2007). A similar argument applies for measurements using the HI Ly α forest power spectrum. The thin curves bracketing the thick ones are ± 1 s.d. of the mean T_0 . Note that the agreement with measurements is better if we compare with the -1 s.d. curves rather than the mean T_0 curve. A more detailed study is required in order to determine whether these observations are consistent with our simulations.

Figure 12 also includes two curves for the case in which HeII reionization happens at $z > 6$ and the simulation is initialized with temperatures of $T(\Delta_b) = 10^4$ K and $T(\Delta_b) = 2 \times 10^4$ K at $z = 6$ (thin cyan dot-dashed curves). By $z \approx 3$ both of these curves have asymptoted to a temperature of $\approx 7,000$ K as expected




FIG. 10.— Power spectrum of x_{HeIII} (top panel), δT (middle panel), and the cross correlation coefficient of x_{HeIII} and δ_b (bottom panel) in simulation B1. The thin blue dotted curve in the middle panel is from a simulation without HeII reionization.




FIG. 11.— PDF of the temperature tabulated from grid cells with $-0.1 < \delta_b < 0.1$ in simulations D2, L1, and S3.




FIG. 12.— Top Panel: Best-fit power-law index of the $T-\Delta_b$ relation in two simulations (thick curves), as well as its evolution in simulations without a $z \sim 3$ HeII reionization (thin curves). The markers with error bars are measurements of this quantity in the literature, and the references for these measurements are given in the key in the bottom panel. Bottom Panel: Evolution of best-fit T_0 in three simulations (thick curves) as well as ± 1 s.d. in this quantity (the thinner curves with the same line style). The thin cyan dot-dashed curves that are increasing with z represent simulations in which HeII reionization occurs at $z > 6$. Also shown are measurements of T_0 .

(Hui & Gnedin 1997). The curves that represent the simulations with HeII reionization (at least in models D1 and L1) are more consistent with the measurement points than these curves.

5.2. Temperature-Density Relation

The temperature of the gas depends primarily on the competition between photo-heating and adiabatic cooling. These processes lead to a power-law relationship between T and Δ_b for unshocked gas at times sufficiently after a reionization epoch (Hui & Gnedin 1997). This power-law index asymptotes in time to the value $\gamma - 1 = 0.6$ (Hui & Gnedin 1997).¹⁹

However, HeII reionization can lead to a more complicated relation between T and Δ_b than a single power-law. Just after an *instantaneous* and *homogeneous* HeII reionization process, the equation of state $\gamma - 1$ would equal zero (an isothermal IGM), at least if the temperature afterward is much larger than before. However, the finite duration of HeII reionization causes $\gamma - 1$ to deviate from zero, and the duration as well as the inhomogeneity of the photo-heating introduces scatter into the $T-\Delta_b$ relationship. Regions that are ionized earliest have more time to cool, and regions ionized by the most filtered radiation are the hottest.

Figure 13 shows the evolution of the average $T-\Delta_b$ relation in simulation D1, which uses Quasar Model II, at times for which $\bar{x}_{\text{HeIII}} = 0, 0.5, 0.8$, and 1 (thick solid curves). The slope of this relation (which was initialized with $\gamma - 1 = 0.3$ at $z = 6$) does not change substantially throughout HeII reionization. While the value of $\gamma - 1$ ap-

¹⁹ This value for $\gamma - 1$ is just slightly smaller than if only adiabatic cooling is included, which would yield $\gamma - 1 = 2/3$.




FIG. 13.— T - Δ_b relation in simulation D1. This simulation is initialized with $\gamma - 1 = 0.3$ at $z = 6.1$. The thick solid curves are the mean $T(\Delta_b)$, and the contours enclose (33, 67, 90, 99, 99.9) percent of the grid cells – encompassing the regions that have the highest densities in $\log T - \log \Delta_b$.

pears fairly constant, the dispersion in this relationship grows during HeII reionization. The contours enclose (33, 67, 90, 99, 99.9) percent of the cells in the simulation grid – encompassing the regions that have the highest densities in $\log T - \log \Delta_b$.

The top panel in Figure 12 shows the measured value of $\gamma - 1$ at various times in simulations L1 and L1b, which use QSO Model II and filtering Method A.²⁰ The only difference between these two simulations is that L1 is initialized with $\gamma - 1 = 0.3$ at $z = 6$ and L1b with $\gamma - 1 = 0$. In both of these cases, as well as in all of the simulations listed in Table 1, the T - Δ_b relation does not become inverted or isothermal (i.e., $\gamma - 1 \leq 0$). This result owes to the large QSO bubbles being essentially uncorrelated with Δ_b ; no value of Δ_b is ionized preferentially at a given time during HeII reionization. Unless some Δ_b is ionized preferentially, then the inequality $\gamma - 1 > 0$ must hold. Our simulations yield $\gamma - 1 \approx 0.3$ during the bulk of HeII reionization. After HeII reionization, this index steepens. Interestingly, the evolution of $\gamma(z) - 1$ in simulations S4, which has $\bar{\alpha}_{UV} = 0.6$, is nearly identical to this in L1, even though the IGM in S4 is heated to a much higher temperature.

Also shown in Figure 12 are measurements of $\gamma - 1$ in the literature. The different estimates are not consistent with one another, and generally have large error bars. These measurements are calibrated using hydrodynamic or hydro-PM simulations that do not account for inhomogeneous heating owing to HeII reionization.²¹

Previous studies have found differing results concern-

²⁰ We fit a power-law to the mean value of $T(\Delta_b)$ in bins spanning $0.5 < \Delta_b < 2$ to obtain $\gamma - 1$. We find that the same procedure applied to the median of $T(\Delta_b)$ rather than the mean or to a more restricted range in Δ_b yields similar results.

²¹ Also, note that the Schaye et al. (2000) data points – which are the most discrepant with respect to our predictions – are each taken from half of the Ly α forest from a *single* quasar spectrum (i.e., approximately a 250 Mpc skewer). Because of the large structures during HeII reionization, it is conceivable that the Schaye et al. (2000) estimates do not sample enough regions to be representative.

ing the effect of HeII reionization on the T - Δ_b relation. As in our simulations, the T - Δ_b relation in the studies of Gleser et al. (2005) and Paschos et al. (2007) is not inverted. In fact, the predictions for $\gamma - 1$ in Gleser et al. (2005) are quite steep with $\gamma - 1 \geq 0.38$. In contrast, Furlanetto & Oh (2007) found that the T - Δ_b relation could become inverted by HeII reionization and that the mean relation could deviate significantly from a power-law in their excursion set based model. This result owed to the underdense regions being ionized later during HeII reionization and, therefore, being the hottest. In a toy case considered in Furlanetto & Oh (2007), in which there are no correlations between the QSO bubble and the density field, the evolution of $\gamma - 1$ during HeII reionization is similar to what we find.

The no-correlation approximation in the Furlanetto & Oh (2007) toy model is reasonable during HeII reionization by quasars. The magnitude of the correlation between HeIII bubbles and fluctuations on the Jeans scale is relevant for how HeII reionization affects the T - Δ_b relation. The linear-theory correlation coefficient between the overdensity in a region of $10^{12} M_\odot$ that hosts a QSO and a Jeans mass region separated from it by 20 Mpc (where we take $M_J = 10^{10} M_\odot$) is $r = 0.02$ ($r = 0.008$ for a 30 Mpc separation). The value of r would be even smaller if non-linearities are taken into account. Therefore, the probability distribution function of δ_b smoothed on the Jeans scale within the HeIII region of a QSO is a Gaussian (in linear theory) centered at $\langle \delta_b \rangle = r N_\sigma$, where N_σ is the quasar overdensity in units of the s.d. in δ_b . The variance of this Gaussian is reduced by $1 - r^2$ from the volume-averaged variance. Therefore, for the small r values quoted above and reasonable values of N_σ , a Jeans mass clump in a quasar HeIII region has a nearly equal chance to be overdense as it does to be underdense.

Bolton et al. (2008a) and other studies have postulated that radiative transfer effects can result in an inverted T - Δ_b relation. The idea is that the radiation which makes it out into the underdense voids is highly filtered, heating these regions to higher temperatures. However, this picture does not apply when quasars are the sources that ionize the HeII. Voids on the Jeans scale, the scale that is most relevant for the Ly α forest transmission, are not strongly anti-correlated with the giant HeIII bubbles. Instead, radiation that is filtered as it traverses a QSO bubble has an approximately equal chance of ionizing a void or filament at the bubble edge. Therefore, the filtering leads to scatter in the relation between T and Δ_b , but not to an inverted relation. If the HeII-ionizing sources were more numerous and much dimmer than QSOs, in this case an inverted relation is possible.

Interestingly, Bolton et al. (2008a) found that the flux probability distribution from the $z \sim 3$ HI Ly α forest favors an inverted T - Δ_b relation to accommodate the measured number of high transmission pixels. We have shown that the observed population of QSOs cannot be responsible for an inverted relation. However, we think that HeII reionization by QSOs could still reconcile the Bolton et al. (2008a) result. The Bolton et al. (2008a) analysis did not consider physically motivated models




FIG. 14.— Evolution of the effective HeII optical depth. The error bars on most of the simulation curves bracket the values in which $\tau_{\text{HeII,eff}}$ falls 90% of the time when measured over a skewer of length 190 Mpc. The measurement points are described in the text.

for the dispersion in this relation.²² The large dispersion in $T(\Delta_b)$ that we find will act to increase the number of unabsorbed pixels and could potentially accommodate the Bolton et al. (2008a) result without the need for an inverted T - Δ_b relation.

5.3. Mean HeII Ly α Forest Transmission

There are currently only a handful of HeII Ly α forest sight-lines. The number of such sight-lines is limited because a bright quasar is required to achieve adequate sensitivity with existing instruments and because a single HI Lyman-limit system can absorb much of the flux at wavelengths where HeII Ly α absorption occurs. The number of HeII Ly α forest sight-lines is expected to increase once the Cosmic Origins Spectrograph is installed on Hubble, which is currently scheduled for spring of 2009.

In this section, we calculate $\tau_{\text{eff, HeII}}$ along skewers through our simulation box. These spectra, as well as those discussed in Section 5.4, are generated by gridding the N-body particles on a 256^3 mesh in the 186 Mpc box and using the temperature and ionization fields from the simulations. Our predictions for the HeII-Ly α transmission are reliable to the extent that our dark matter simulations capture the gas density PDF (especially for $\delta_b < 0$, where most of the transmission occurs). In Appendix D, we show that the dark matter PDF measured from our grid reasonably reproduces the expected gas PDF. Our predictions for the mean HeII Ly α transmission change by only $\lesssim 20\%$ when we calculate this statistic from the 512^3 gridded N-body field rather than the 256^3 one. This agreement is reassuring and suggests that our predictions are robust.

Figure 14 plots $\tau_{\text{eff, HeII}}$ from our simulations and from the most constraining observations, where $\tau_{\text{eff, HeII}}$

is defined as minus the log of the HeII Ly α forest transmission. The square markers are from observations of quasars HS1700+6416 (Fechner et al. 2006), HS1157+3143 (Reimers et al. 2005), PKS1935-692 (Anderson et al. 1999), Q0302-003 (Heap et al. 2000), and SDSSJ2346-0016 (Zheng et al. 2004). The vertical error bars on these measurements are just statistical, and the horizontal error bars signify the redshift range over which each of the values were estimated.²³

The curves in Figure 14 are $\tau_{\text{eff, HeII}}$ from four simulations, for which each marker value is estimated from 1000, 186 Mpc random skewers through the simulation volume. The error bars on most of the simulation curves bracket the values in which $\tau_{\text{eff, HeII}}$ falls 90% of the time when measured along a single skewer. Note that the measurement points were taken from skewers that are comparable in length to the 186 Mpc box. There should be an additional cosmic variance error on these points that is similar in magnitude to the error bars on the simulation curves.

Our simulated predictions for $\tau_{\text{eff, HeII}}(z)$ provide moderate agreement with the measured values. The evolution of $\tau_{\text{eff, HeII}}(z)$ in simulations D1 and L1c is more consistent with three of the four measured points at $z < 3$ than the $\tau_{\text{eff, HeII}}(z)$ in simulation L1. None of the simulations are consistent with the $z = 2.7$ measurement (although, simulation D1 was terminated at $z = 2.9$ and would have been the most consistent).

The differences in $\tau_{\text{eff, HeII}}(z)$ among the simulations stem in part from reionization ending at $z = 3.15$ in D1 and $z = 3.45$ in L1c, whereas it is completed at a later time, at $z = 2.95$, in simulations L1 and L3. (The end of HeII reionization is defined here as when $\bar{x}_{\text{HeIII}}(z_{\text{end}}) = 0.95$.) Another difference is that, at fixed \bar{x}_{HeIII} , the value of $\bar{\Gamma}_{\text{HeII}}$ is smaller in simulations L1, L1c, and L3 than it is in D1 because the former supplement the resolution by using filtering Method A or B while the latter does not employ additional filtering. A larger value for $\bar{\Gamma}_{\text{HeII}}$ results in more transmission.

Simulations L1 and L3 differ only in the filtering method that is used, and this also leads to differences in the predictions for $\tau_{\text{eff, HeII}}(z)$ primarily after HeII reionization. The uncertainties in the filtering of absorption systems can be remedied with improvements to the radiative transfer algorithm and by using high resolution gas simulations. Even though some theoretical uncertainty remains, the different trends among the simulations suggest that $\tau_{\text{eff, HeII}}(z)$ could be useful to distinguish between different HeII reionization models.

5.4. Mean HI Ly α Forest Transmission

By increasing the temperature of the IGM, HeII reionization changes the photo-ionization state of the intergalactic hydrogen (because $\bar{x}_{\text{HI}} \approx \alpha_A n_e / \Gamma_{\text{HI}}$, where n_e is the electron number density and $\alpha_A \sim T^{-0.7}$), thereby affecting the transmission properties of the HI Ly α forest. Interestingly, Bernardi et al. (2003) found evidence for a depression in the mean transmission at $z \approx 3.2$ with width $\Delta z \approx 0.3$. A similar depression was later

²² They did consider a model with $\gamma - 1 \approx 0.3$ in which this scatter was increased artificially by adding a Gaussian dispersion to the particle photo-heating rates, and they found that this model resulted in essentially no change in the flux PDF over a model without scatter. However, the Bolton et al. (2008a) model with scatter resulted in similar 90% contours in the T - Δ_b plane compared to their model without scatter (see their Fig. 3 and compare with our Fig. 13). The width of their 90% contours in the T direction are factors of several smaller than what we find.

²³ Often the span over which a measurement of $\tau_{\text{eff, HeII}}$ is done is chosen such that it encompasses a very dark (or bright) region in the HeII Ly α forest. Unfortunately, this practice significantly biases the observational estimates.




FIG. 15.— Evolution of the effective HI Ly α forest optical depth. The red solid points with error bars are the measurement of $\tau_{\text{eff, HI}}$ from Faucher-Giguère et al. (2008a), and the black crosses are the estimates of Bernardi et al. (2003). The features that are present at $z \approx 3.2$ in the Bernardi et al. (2003) and Faucher-Giguère et al. (2008a) data cannot be explained by HeII reionization in our simulations. Any disagreement between the data points and the model curves does not imply that the model is disfavored, and instead it may indicate that an incorrect form for $\Gamma_{\text{HI}}(z)$ was assumed.

discovered by Faucher-Giguère et al. (2008a).²⁴

The red points with solid error bars in Figure 15 are the Faucher-Giguère et al. (2008a) measurement of $\tau_{\text{eff, HI}}$, which is defined as minus the log of the HI Ly α forest transmission. The black crosses with errors are the Bernardi et al. (2003) measurement. All errors bracket the 68% confidence level region. Neighboring points in the Bernardi et al. (2003) measurement are correlated. The dip in absorption in these measurements that has been the subject of attention is at $z \approx 3.2$ and has width $\Delta z \lesssim 0.4$. A dip is present with higher significance in the Bernardi et al. (2003) measurement.

The theoretical curves are calculated from our simulations in the manner described in Section 5.3 except for HI Ly α absorption. These calculations assume a uniform HI-ionizing background with $\Gamma_{\text{HI}}(z) = 10^{-12} \text{ s}^{-1}$. The effective optical depth in simulation L1 is representative of most of our runs, and $\tau_{\text{eff, HI}}$ in simulation S4b is smaller than is typical because of the higher IGM temperatures achieved in this simulation. The dashed curves give the evolution of the transmission for a simulation that is initialized with $\gamma - 1 = 0$ and $T_0 = 10^4 \text{ K}$ or $2 \times 10^4 \text{ K}$ at $z = 6$ and in which HeII reionization does not happen at $z < 6$. Any disagreement between the data points and the model curves does not imply that the model is disfavored. The form of $\Gamma_{\text{HI}}(z)$ can be adjusted to yield any form for $\tau_{\text{eff, HI}}(z)$.

Our predictions for $\tau_{\text{eff, HI}}(z)$ show no feature owing to HeII reionization. HeII reionization takes place over a much longer interval in the simulations than the $\Delta z \lesssim 0.4$ of the features seen in these measurements. The duration of HeII reionization in our simulations is determined primarily by the density of $L \sim L_*$ quasars, which over the relevant redshift interval is well constrained by current observations. In addition, even once HeII reioniza-

tion is complete, it takes $\Delta z \sim 1.5$ after HeII reionization for the gas at mean density to cool by a factor of two. Therefore, cooling after HeII reionization is unlikely to produce any sharp feature in $\tau_{\text{eff, HI}}$.

Theuns et al. (2002b) found using hydrodynamic simulations that $\tau_{\text{eff, HI}}$ can evolve significantly over a shorter interval than the time over which the gas cools. They explained this effect as owing to velocity gradients that are induced by the heating. They claimed that these gradients broaden absorption lines, creating the rise in $\tau_{\text{eff, HI}}$. Furthermore, they found that the rise after HeII reionization in their simulations is similar to the rise seen at the low redshift end of the Bernardi et al. (2003) feature. This effect is not captured by our method, which misses the backreaction of the photo-heating on the gas density. However, Theuns et al. (2002b) attributed the decrease in $\tau_{\text{eff, HI}}$ on the other side of the Bernardi et al. (2003) feature to heating from HeII reionization, which they modeled as a quick, homogeneous process. In our simulations, HeII reionization takes place over $\Delta z > 1$ compared to $\Delta z \approx 0.1$ in the Theuns et al. (2002b) simulations, and, therefore, any trends in $\tau_{\text{eff, HI}}$ owing to heating must be broader in our picture.

6. CONCLUSIONS

We have run a set of simulations of HeII reionization to understand the structure of HeII reionization and its effect on several observables. We find that for a late reionization of HeII by quasars:

- The popular assumptions that the HeIII ionization fronts are sharp and that there is uniform heating within the front (and no heating outside of it) do not yield a realistic model for HeII reionization. While the ionization fronts are still fairly localized, hard photons stream far from their sources and are absorbed ahead of the front. These photons inject at large distances a significant fraction of the energy radiated above the HeII Lyman-limit.
- The average temperature at the mean density is increased by $\approx 12,000 \text{ K}$ over the average temperature of the gas in the absence of HeII reionization for $\bar{\alpha}_{\text{UV}} = 1.6$. This temperature increase is consistent with estimates that assume that the IGM absorbs all photons with energies less than a few hundred eV during HeII reionization. Regions that are ionized last are ionized by the hardest radiation, reaching $T > 30,000 \text{ K}$ in our fiducial model. If the spectral index of the QSOs is different from our fiducial value of $\bar{\alpha}_{\text{UV}} = 1.6$, the average temperature can be significantly altered.
- Poisson fluctuations in the number of QSOs rather than their spatial clustering shape the structure of the ionization and temperature fluctuations on $\lesssim 50 \text{ Mpc}$ scales because rare $L \sim L_*$ quasars ionize the HeII. HeII reionization produces large temperature fluctuations on 50 Mpc scales, and it results in the ionization and HeII photo-heating fluctuations being essentially uncorrelated with the density fluctuations on the Jeans scale, the scale most relevant for HI and HeII Ly α forest statistics.
- Measurements of the $z \sim 3$ forest suggest $T_0 \approx 20,000 \text{ K}$. Without invoking an exotic heating

²⁴ The Faucher-Giguère et al. (2008a) estimate uses 84 high resolution Keck and Magellan spectra, whereas the Bernardi et al. (2003) measurement is from 1061 lower signal to noise and lower resolution Sloan spectra.

mechanism, a late HeII reionization epoch is required to produce this temperature. The amount of additional heat provided by HeII reionization in our simulations is enough to produce the inferred temperatures.

- To the extent that there is a T - Δ_b relation, HeII reionization by quasars leads to a temperature-density equation of state of $\gamma - 1 \approx 0.3$ for $0.1 \lesssim \bar{x}_{\text{HeIII}} \lesssim 0.9$. HeII reionization by QSOs cannot result in an inverted relation ($\gamma - 1 < 0$) as has been claimed.
- Our simulations of a $z \sim 3$ HeII reionization process produce a similar evolution in the HeII Ly α mean transmission to what has been estimated. Better observations of the mean transmission and its scatter will be able to rule out models presented here.
- In our simulations, the heating from HeII reionization by quasars is unable to produce any semblance of the observed depression in the $z \approx 3.2$ HI Ly α forest opacity. The HeII reionization process is too extended in redshift ($\Delta z > 1$) to be responsible for this narrow feature ($\Delta z \lesssim 0.4$).

The morphology of HeII reionization is considerably different than that of hydrogen reionization (e.g., McQuinn et al. 2007). During hydrogen reionization, current models find that the growth of ionized bubbles is more collective; hundreds or even thousands of galaxies within a bubble contribute to its growth, leading to the bubble structure tracing the distribution of galaxies and to tens of Mpc HII regions (e.g., Zahn et al. 2007). During HeII reionization, the growth is more stochastic. Regions that happen to host a “nearby” quasar (within ~ 30 Mpc) are ionized by that quasar. Because the QSO bubbles are so extended, the ionized structures are larger than during HI reionization.

Another significant difference stems from the m.f.p. of the ionizing photons to be absorbed in diffuse gas. The

spectrum from quasars is harder than that of stars – our best guess for what ionizes the hydrogen – the number density of helium is 70 times smaller at $z = 3$ than hydrogen at $z = 6$, and the cross section is 4 times smaller. These three factors conspire to make the typical m.f.p. for a HeII ionizing photon megaparsecs rather than kiloparsecs, as it is for HI ionizing photons during hydrogen reionization. We have seen that some HeII ionizing photons free stream far from their sources, partially ionizing and heating up these regions. If stars reionize the hydrogen, ionization and heating occurs within HII bubbles.

A definitive identification of the redshifts of HeII reionization will place constraints on the sources that produce > 4 Ry photons and will aid studies of the HI Ly α forest. The data from previous observations, if analyzed properly, may be able to confirm whether quasars reionize HeII at $z \approx 3$. In addition, future observations of the HI and HeII Ly α forest will soon be available with Sloan Digital Sky Survey III²⁵ and the Cosmic Origins Spectrograph on the Hubble Space Telescope. These telescopes will significantly increase the number of sight-lines in the HI and HeII Ly α forests. It is timely to make predictions for the effect of HeII reionization on these observations.

7. ACKNOWLEDGMENTS

Many thanks to Mark Dijkstra, Alexandre Tchekhovskoi, and Hy Trac for interesting discussions, and to James Bolton and Steven Furlanetto for providing useful comments on the manuscript. MM is supported by the NSF graduate student fellowship. This work was supported in part by NSF grants ACI 96-19019, AST 00-71019, AST 02-06299, AST 03-07690, and AST 05-06556, and NASA ATP grants NAG5-12140, NAG5-13292, NAG5-13381, and NNG-05GJ40G. Further support was provided by the David and Lucile Packard, the Alfred P. Sloan, and the John D. and Catherine T. MacArthur Foundations.

²⁵ www.sdss3.org

APPENDIX

A. HEII ABSORBERS

In this section, we examine the effect that dense systems have on filtering the radiation field and whether this effect is adequately captured in our simulations. We discuss two methods that we use to supplement the resolution of the simulations.

Analytic Modeling

The HI column density (N_{HI}) distribution is relatively well measured from the HI Ly α forest, whereas the HeII column density (N_{HeII}) distribution is significantly more uncertain. Given the well-measured N_{HI} distribution, we would like to derive the N_{HeII} distribution and then use this to calculate the effect of filtering. We follow the method described in Haardt & Madau (1996). This approach uses the value of N_{HI} for an absorber and the incident values of Γ_{HI} and Γ_{HeII} to predict the value of N_{HeII} .

If the absorber is optically thin to 4 Ry photons, it is trivial to infer the value of N_{HeII} , namely

$$N_{\text{HeII}} = N_{\text{HI}} \eta_{\text{thin}}, \quad (\text{A1})$$

where

$$\eta_{\text{thin}} \approx \frac{5.5 Y_{\text{He}}}{4 (1 - Y_{\text{He}})} \left(\frac{\Gamma_{\text{HI}}}{\Gamma_{\text{HeII}}} \right), \quad (\text{A2})$$

and where Y_{He} is the mass fraction in helium and the factor 5.5 is the approximate ratio of HeII and HI recombination rates. HI column densities with

$$N_{\text{HI,thin}} < 4 \times \sigma_{\text{HI}}^{-1} \eta^{-1} = 1.3 \times 10^{16} \left(\frac{50}{\eta} \right) \text{ cm}^{-2}, \quad (\text{A3})$$

yield optical depths of $\tau \lesssim 1$ for 4 Ry photons, where $\eta \equiv N_{\text{HeII}}/N_{\text{HI}}$.

What values of η do we expect? Taking the Γ -values inferred from the HI Ly α and HeII Ly α forests at $z \approx 2.5$ of $\Gamma_{\text{HI}} \approx 10^{-12} \text{ s}^{-1}$ and $\Gamma_{\text{HeII}} \approx 10^{-14} \text{ s}^{-1}$ (e.g., Shull et al. 2004), equation (A2) yields $\eta_{\text{thin}} \approx 50$. Or, taking the $z \approx 3$ values of $\Gamma_{\text{HI}} \approx 10^{-12} \text{ s}^{-1}$ and $\Gamma_{\text{HeII}} \approx 10^{-15} \text{ s}^{-1}$ – at least in some regions – η_{thin} increases to 500 (e.g., Reimers et al. 1997).

An absorber is optically thick to 4 Ry photons while also being optically thin to 1 Ry photons for $N_{\text{HI,thin}} \lesssim N_{\text{HI}} \lesssim \sigma_{\text{HI}}^{-1} = 1.6 \times 10^{17} \text{ cm}^{-2}$. Consequently, Γ_{HeII} can be much lower within an absorber than the incident value, and, therefore, N_{HeII} can be larger relative to N_{HI} than in the optically thin case. Thus, to calculate N_{HeII} we solve the quadratic formula given in Fardal et al. (1998) (their eqn. A11), which determines N_{HeII} assuming a slab of primordial gas. This equation requires as input N_{HI} , Γ_{HI} , Γ_{HeII} , and n_H – the number density of atomic and ionized hydrogen in the absorber. This method is an improvement over the Haardt & Madau (1996) multi-zoned model for incorporating HeII absorption (Fardal et al. 1998).

To proceed with solving the Fardal et al. (1998) equation for N_{HeII} , a model for $n_H(N_{\text{HI}})$ is required. Simulations of the HI Ly α forest have had remarkable success at matching the N_{HI} distribution found in HI Ly α forest observations (e.g., Cen et al. 1994; Miralda-Escudé et al. 1996; Hernquist et al. 1996; Katz et al. 1996), and analytic models have been constructed to understand these simulations. In particular, Schaye (2001) argued that Ly α forest absorbers have sizes that are comparable to the Jeans length, $L_J(\delta_b, T)$. With the assumption that $N_{\text{HI}} = n_{\text{HI}} L_J$ and of photo-ionization equilibrium, it is trivial to map between n_H and N_{HI} for $N_{\text{HI}} \lesssim \sigma_{\text{HI}}^{-1}$, namely (Schaye 2001)

$$n_H \approx 6 \times 10^{-4} \text{ cm}^3 \left(\frac{N_{\text{HI}}}{10^{16} \text{ cm}^2} \right)^{2/3} T_4^{0.17} \left(\frac{\Gamma_{\text{HI}}}{10^{-12} \text{ s}^{-1}} \right)^{2/3}, \quad (\text{A4})$$

where $T_4 = T/10^4 \text{ K}$. This formula agrees well with simulations (Schaye 2001). For our purposes it may reasonable to use equation A4 even when $N_{\text{HI}} \gtrsim \sigma_{\text{HI}}^{-1}$, which is beyond its realm of applicability, because such high column density absorbers play an insignificant role in filtering the radiation near the HeII Lyman-limit.

Figure 16 plots the predictions of this model for $\eta(N_{\text{HI}})$ (bottom panel), the m.f.p. to be absorbed in systems with column density greater than N_{HI} (top panel), and the contribution of different systems to the m.f.p. ($d\lambda_{\text{mfp}}(> N_{\text{HI}})^{-1}/d\log N_{\text{HI}}$ (middle panel). The function $d\lambda_{\text{mfp}}(> N_{\text{HI}})^{-1}/dN_{\text{HI}}$ is proportional to the probability a photon is absorbed in a system of column density N_{HI} . These panels use the A1 fit to the N_{HI} distribution in Fardal et al. (1998).²⁶ All curves assume $\Gamma_{\text{HI}} = 10^{-12} \text{ s}^{-1}$ and $T = 10^4 \text{ K}$. In the top panel, the solid curves take $\Gamma_{\text{HeII}} = 10^{-15} \text{ s}^{-1}$, and the dotted curves take $\Gamma_{\text{HeII}} = 10^{-14} \text{ s}^{-1}$. The vertical lines in both panels correspond to the approximate value of N_{HI} at which an absorber becomes self-shielded to HeII ionizing photons for $\Gamma_{\text{HeII}} = 10^{-14} \text{ s}^{-1}$ (solid line) and $\Gamma_{\text{HeII}} = 10^{-15} \text{ s}^{-1}$ (dotted line). For the column densities at which an absorber becomes self-shielded, the value of η increases. However, this increase is compensated by the monotonically increasing nature of $n_H(N_{\text{HI}})$, and ultimately this effect wins out and $\eta(N_{\text{HI}})$ decreases. This decrease is important, allowing 100 eV photons to have long m.f.p.s to be absorbed in dense systems.

For $\Gamma_{\text{HeII}} = 10^{-14} \text{ s}^{-1}$, absorbers with $N_{\text{HI}} \gtrsim 10^{16} \text{ cm}^{-2}$ limit the m.f.p. of 4 Ry photons to $\approx 40 \text{ Mpc}$ at the HeII Lyman-limit (see thickest, solid curve in middle panel in Fig. 16). For $\Gamma_{\text{HeII}} = 10^{-15} \text{ s}^{-1}$, lower column density absorbers with $10^{14} \gtrsim N_{\text{HI}} \gtrsim 10^{15} \text{ cm}^{-2}$ limit the m.f.p. of these photons to $\approx 10 \text{ Mpc}$ (dashed curves in middle and top panels). Higher energy photons have longer m.f.p. values.

Filtering in our Simulations

The radiative transfer grid employed in our study may not capture all of the absorbers with $N_{\text{HI}} \gtrsim 10^{15} \text{ cm}^{-2}$ (Appendix D). These absorbers can play an important role in filtering the radiation field. Furthermore, even if these absorbers are resolved on the grid, the code will not properly handle the filtering when $\Gamma_{\text{HeII}} \gtrsim \Delta t^{-1}$. We describe below two methods that we use to better capture these effects.

Filtering Method A

The simpler of the two filtering methods uses the density structure on the radiative transfer grid and a modification of the radiative transfer code to achieve the appropriate filtering. Appendix D shows that the column density distribution of absorbers measured from the 256^3 and 512^3 N-body simulation grids reproduce to within a factor of 2 the column density distribution seen in gas simulations. Therefore, our gridded N-body simulation will produce the correct filtering of the radiation field if it passes three tests:

²⁶ Fardal et al. (1998) derived best fit functional forms for the distribution of absorbers as a function of N_{HI} and z (see Table 1 in Fardal et al. 1998), fitting to a compilation of Ly α forest observations. These fits account for the observed depression in $\partial^2 N / \partial N_{\text{HI}} \partial z$ centered at $N_{\text{HI}} \approx 10^{16} \text{ cm}^{-2}$ (Petitjean et al. 1993). Note that the piecewise form for the curves in the middle panel in Figure 16 owes to the piecewise power-law that Fardal et al. (1998) used to fit the observed N_{HI} distribution.




FIG. 16.— Dotted curves in each panel are for $\Gamma_{\text{HeII}} = 10^{-15} \text{ s}^{-1}$ and the solid are for $\Gamma_{\text{HeII}} = 10^{-14} \text{ s}^{-1}$ and all curves use $\Gamma_{\text{HI}} = 10^{-12} \text{ s}^{-1}$. Top Panel: The m.f.p. for photons at select energies to be absorbed in a system with column density greater than N_{HI} . Middle Panel: The derivative with respect to N_{HI} of the inverse of the m.f.p. This quantity is maximized for the N_{HI} that contribute most significantly to the absorption. Bottom Panel: The value of $\eta = N_{\text{HeII}}/N_{\text{HI}}$ in this model as a function of N_{HI} . The vertical lines in both panels correspond to the N_{HI} at which an absorber becomes optically thick for the Γ_{HeII} with the same line style.

- The density structure of the absorbers is captured to the level that is required.
- The radiative transfer method accounts for the filtering of the radiation field as radiation traverses overdense cells.
- The clustering of the absorbers is taken into account.

When the absorber is optically thin, η is independent of the absorber density. As seen in the middle panel in Figure 16, a significant fraction of the filtering owes to optically thin absorbers when $\Gamma_{\text{HeII}} = 10^{-15} \text{ s}^{-1}$. Even when an absorber becomes optically thick, the dependence of η on density is weak for the most relevant column densities if one uses the Schaye (2001) model for $n_H(N_{\text{HI}})$. (Note the flatness of η in the bottom panel in Fig. 16 for the N_{HI} that contribute the most to the m.f.p.) Therefore, since the N_{HI} column density distribution is captured fairly well at our grid scale, there is reason to believe that the density is adequately described for the purpose of filtering the radiation field.

Passing the second of the three tests requires an addition to the radiative transfer algorithm beyond what was outlined in the previous section. For $\Gamma_{\text{HeII}}^{-1} \ll \Delta t$, a cell in the code outlined in §2 would have $x_{\text{HeII}} < x_{\text{HeII,eq}}$ once some fraction of the rays that will travel through that cell in a timestep have done so. Therefore, the filtering will be underestimated for the rays that enter after $x_{\text{HeII}} < x_{\text{HeII,eq}}$. To remedy this, once a cell with $\Delta_b > \Delta_F$ has $x_{\text{HeII}} < x_{\text{HeII,eq}}$, we perform the radiative transfer as if $x_{\text{HeII}} = \tilde{x}_{\text{HeII,eq}}$, where $\tilde{x}_{\text{HeII,eq}} = \alpha_A n_e / \tilde{\Gamma}_{\text{HeII}}$ and $\tilde{\Gamma}_{\text{HeII}}$ is tabulated from all the previous rays that have entered the cell within the timestep up to the current ray. Note that x_{HeII} always approaches $x_{\text{HeII,eq}}$ from below because recombinations are performed prior to ionizations within a

timestep. We set $\Delta_F = 4$, which for $\Gamma_{\text{HI}} = 10^{-12} \text{ s}^{-1}$ corresponds to

$$N_{\text{HI}} = 2 \times 10^{14} \text{ cm}^{-2} \left(\frac{\Delta_F}{4} \right)^2 \left(\frac{l_{\text{cell}}}{1 \text{ Mpc}} \right) \left(\frac{1+z}{4} \right)^5. \quad (\text{A5})$$

Passing the final of the three tests – capturing the clustering of the absorbers – is the most difficult and is not achieved with this method. Even though the column density distribution is maintained, the bias of the absorbers is changed by the coarse gridding (Appendix D). It is conceivable that capturing the bias is not of crucial importance because the m.f.p. for ionizing photons is typically larger than the correlation length of the absorbing systems. This method does not filter the background radiation ($\gtrsim 500 \text{ eV}$). Figure 16 suggests that such filtering is not important because the m.f.p. for these photons to be absorbed in a dense system is gigaparsecs.

Filtering Method B

The second method for filtering the radiation is based on the model proposed in Haardt & Madau (1996) and improved in Fardal et al. (1998) to infer N_{HeII} from N_{HI} , which is described in the beginning of this section. This inference also requires as input Γ_{HeII} . We take it to be the maximum of the current tabulated value of Γ_{HeII} in the cell and the value of Γ_{HeII} from the previous timestep.

In addition, this method assumes that Ly α absorption arises from gas at the outskirts of dark matter halos. We assume that the cross section for a photon to intersect an absorber of column density N_{HeII} associated with a halo of mass m has the functional form

$$\sigma(m, N_{\text{HeII}}) = \sigma_{m_*} P(N_{\text{HeII}}) (m/m_*)^{2/3}, \quad (\text{A6})$$

where $P(N_{\text{HeII}})$ is taken to be the probability distribution of N_{HeII} that corresponds to lines with $\log_{10}(N_{\text{HI}}) > 14.5$ and is zero for smaller column densities, and σ_{m_*} sets the normalization. In this method, rays that have photons of energy E_γ and that travel through a cell with a halo of mass m are attenuated by the factor $\int dN_{\text{HeII}} \sigma(m, N_{\text{HeII}}) \exp[-\tau(E_\gamma, N_{\text{HeII}}, N_{\text{HeII}})] / l_{\text{cell}}^2$, where $\tau = \sigma_{\text{HI}}(E_\gamma) N_{\text{HI}} + \sigma_{\text{HeII}}(E_\gamma) N_{\text{HeII}}$.

The $m^{2/3}$ scaling in equation (A6) is chosen such that the cross section is proportional to the square of the halo virial radius. The values of σ_{m_*} and m_* are assigned so that the mass integral over $\sigma(m, N_{\text{HeII}}) n_h(m)$, where $n_h(m)$ is the halo mass function, yields $dN/dldN_{\text{HeII}}$ – the number of absorbers per proper length l per N_{HeII} . To derive $dN/dldN_{\text{HeII}}$, we use the A1 fit to $dN/dldN_{\text{HI}}(z)$ provided in Fardal et al. (1998) and that $dN/dldN_{\text{HeII}}(z) \equiv \eta(N_{\text{HI}})^{-1} dN/dldN_{\text{HI}}(z)$. While we use the observations to construct $P(N_{\text{HeII}})$, physically it has to do with the density profile around halos. The motivation for this filtering method is that the high column density absorbers are shown to be associated with halos (Cen & Simcoe 1997; Miralda-Escudé 1998). This method filters both the rays and the background radiation.

B. TESTS OF CODE

Number of Photon Bins

Since the computation time scales linearly with the number of frequency bins n_ν and since memory requirements also increase linearly with n_ν , it is important to find the minimum value of n_ν that achieves acceptable convergence. To test how our code converges with the number of frequency bins, we place a quasar in a homogeneous IGM initialized with $x_{\text{HeII}} = 1$ and solve for x_{HeII} and T while varying n_ν . Figure 17 displays the temperature (top panel) and x_{HeIII} (bottom panel) for $n_\nu = 3, 5$, and 50. For simplicity, we have set the speed of light to infinity. The left-most curves are for a source with $\alpha_{\text{UV}} = 1.5$ and $\dot{N} = 10^{54}$ ionizing photons s^{-1} after $t = 50 \text{ Myr}$ and the right-most curves are the same after $t = 250 \text{ Myr}$, where $t = 0$ corresponds to $z = 4$. The curves at fixed t have only minor differences on scales that have $x_{\text{HeII}} \gtrsim 0.01$. However, for $x_{\text{HeIII}} < 0.01$, 3 frequency bins provides fairly poor convergence to the true solution (which essentially is represented by the $n_\nu = 50$ case), whereas 5 frequency bins shows much better convergence. The cosmological simulations presented in this paper use $n_\nu = 5$ unless otherwise specified.

Even though for a single source the solution is well converged to fairly large radii for $n_\nu = 5$, in the simulations the small inaccuracies at large radii add up and can result in a significant error. Figure 18 compares at $\bar{x}_{\text{HeIII}} = 0.85$ of the fiducial 5 frequency bin simulation – simulation D1 – to one with 10 frequency bins – D3 – with the bins centered at (57, 66, 76, 90, 107, 129, 162, 209, 283, 409) eV such that each bin carries an equal energy. We have checked that the latter simulation is well converged to the full solution by running simulations between 5 and 10 bins and with different spacings in energy. The ionization field is similar between simulations D1 and D3, with the voids being slightly more ionized in D1. The value of \bar{T} is 300 K greater in D1, and it can be seen that some of the hotter regions are a bit hotter in D1 than in D3. However, these differences are small compared to the differences that arise owing to, for example, the uncertainty in $\bar{\alpha}_{\text{UV}}$. We have also compared the temperature PDF between D1 and D3, and the PDFs are almost identical. The temperature PDF still extends to 35,000 K in D3 as in D1 at the end of HeII reionization.

Code Comparison

Analytic solutions to multi-frequency radiative transfer problems generally do not exist. To test our code, we compare it with the 1-D radiative transfer code presented in Lidz et al. (2007). To simplify the comparison, we use a 1-D version of our code that is identical to our 3-D code except that it sends a single ray during a timestep to calculate the spherically symmetric solution. The Lidz et al. (2007) code was written independently by Adam Lidz who was




FIG. 17.— Test of convergence with the number of frequency channels for the case of a single QSO in a homogeneous IGM with $x_{\text{HeII}} = 1$. The top panel shows the temperature, and the bottom shows the HeIII fraction. The left-most curves are for $t = 50$ Myr after a source with $\alpha_{\text{UV}} = 1.5$ and $\dot{N} = 10^{54}$ ionizing photons s^{-1} turns on and the right-most curves are $t = 250$ Myr after.




FIG. 18.— Comparison of the fiducial 5 frequency bin simulation – simulation D1 – to a simulation with 10 frequency bins – D3. The panels are taken from the $z = 3.3$ snapshot, for which $\bar{x}_{\text{HeIII}} = 0.85$.

not involved in the development of our new code. The methodology used in our code is significantly different from that used in the Lidz et al. (2007) code. The latter code assumes a power-law spectrum, and uses this power-law and the intervening column densities of HI, HeI, and HeII to infer the incident spectrum at some radius from the source. Furthermore, it uses much smaller timesteps than our code and iterates to achieve convergence.

We analyze the solutions provided by these codes for the simple case of a quasar that produces 10^{54} s^{-1} HeII ionizing photons with a spectral index of 1.5 in a homogeneous IGM with $x_{\text{HeII}} = 1$. We solve for the ionization front of this quasar assuming a homogeneous, static (non-expanding) universe, with density equal to the mean density at $z = 4$. Initially, the hydrogen is assumed to be ionized and the helium to be singly ionized. For this calculation, the speed of light is taken to be infinite, the code used in this paper is set to have 1 Myr timesteps, and the same ray-casting parameters as used in the cosmological runs presented in this paper.

Figure 19 compares the evolution of the temperature and \bar{x}_{HeIII} front as a function of time between a 1-D version of our code (dotted curves) and the code in Lidz et al. (2007) (solid curves). The left-most curves are the temperature, T , and \bar{x}_{HeIII} values as a function of radius after 2 Myr (2 timesteps for our code), the middle curves are those after 8 Myr (8 timesteps), and the right-most curves are those after 20 Myr (20 timesteps). At $r < 5$ Mpc, the agreement between the $T(r)$ produced from the two codes is not perfect (top panel). The slight disagreement owes to a simplification our




FIG. 19.— Comparison of the temperature (top panel) and HeII ionized fraction (bottom panel) between our code (dotted curves) and the code in Lidz et al. (2007) (solid curves). The QSO is shining in a homogeneous medium at $z = 4$ with $\dot{N} = 10^{54}$ HeII ionizing photons s^{-1} . The left-most curves are $T(r)$ and $\bar{x}_{\text{HeIII}}(r)$ after 2 Myr, the middle curves are these functions after 8 Myr, and the right-most curves are after 20 Myr.

algorithm makes: it groups all of the photons released in a timestep into a single photon bundle. In reality, harder photons will travel ahead of softer photons because softer photons are absorbed first. We have verified that if we decrease our timestep the solution from our code converges to that of the Lidz et al. (2007) code. However, our 3-D code sends multiple rays for each cell at the radii where the discrepancy exists, capturing better the region where the 1-D code has not converged. We have verified that this minor discrepancy essentially disappears in the 3-D code. Furthermore, even though timesteps are larger in the cosmological runs than in this test, the convergence of this code is most dependent on the number of timesteps rather than the timestep duration or source luminosity.

The bottom panel in Figure 19 compares the HeIII ionization fronts in the two codes. We plot the y-axis in log because the solutions for the fronts are essentially indistinguishable on linear scales. At large radii, there is a minor systematic difference in the ionization level. This minor difference is much smaller than other uncertainties inherent in HeII reionization calculations, in particular the value of α_{UV} . This systematic does not depend on the number of frequency bins (for $n_\nu \geq 5$) or decreasing the timestep in our code. Also, note that the output of the Lidz et al. (2007) curve is at 1.9 Myr rather than 2 Myr for its left-most curve, which accounts for some of the discrepancy. The maximum differences between the 8 Myr curves is 3%. The small differences at large r could arise because the Lidz et al. (2007) code follow the full three level system, whereas our code ignores HeI.

Temperature-Density Relation

Here, we investigate how well our scheme for computing the temperature evolution in N-body simulations works. The most difficult aspect of such a scheme is to evaluate the convective derivatives that account for the flow of matter through cells. Our scheme for capturing this is discussed in Section 2.2. To test our code, we compare the resulting $T - \Delta_b$ relation in a simulation without HeII reionization to that of the Hui & Gnedin (1997) analytic model. This analytic model has been shown to reasonably match the relation found in semi-analytic calculations that use the Zel'dovich approximation and the evolution of the temperature density relation found in hydrodynamic simulations (although, the Hui & Gnedin (1997) comparison is rather limited for the latter case).

Figure 20 compares the evolution of the $T - \Delta_b$ equation of state $\gamma - 1$ in our code using the 256^3 grid in the 186 Mpc box to the $\gamma - 1$ predicted by the analytic formula in Hui & Gnedin (1997) for models where $z_{\text{rei}} = 6.2$. Our $\gamma - 1$ are calculated by fitting the $T - \Delta_b$ relation to a power-law, weighting the fit by the number of grid cells at each Δ_b . We have checked that our results do not change if we instead fit only to Δ_b values near unity, which is closer to what Hui & Gnedin (1997) calculate.

The agreement is acceptable between the two predictions for $\gamma - 1$. We would not expect exact agreement. Our code calculates the temperature evolution on a non-linear scale (the cell scale) whereas the Hui & Gnedin (1997) formula uses linear theory. Therefore, our value for $\gamma - 1$ should evolve more quickly than the Hui & Gnedin (1997) estimate since Δ_b on average grows faster on non-linear scales.

Our scheme misses the effects of shocks, which add dispersion to the $T - \Delta_b$ relation. The s.d. in the $T - \Delta_b$ relation from shocks is $\lesssim 500$ K at $\Delta_b = 0.5$ and $\lesssim 1000$ K at $\Delta_b = 1$, and it increases steadily as Δ_b increases (see Figure 1 in Hui & Gnedin 1997). The dispersion in temperature at low densities in simulations that include shocks (and a uniform ionizing background) is much smaller than the dispersion that results from HeII reionization.

C. IONIZATION CORRECTION

The fraction of helium in HeII evolves according to the equation (assuming HeI is in ionization equilibrium):

$$\frac{dx_{\text{HeII}}}{dt} = -\Gamma_{\text{HeII}} x_{\text{HeII}} + \alpha_A n_e x_{\text{HeIII}}. \quad (\text{C1})$$




FIG. 20.— Evolution of $\gamma - 1$ in our code compared to that predicted by Hui & Gnedin (1997) (HG97) for $T_{\text{rei}} = 20,000$ K and $T_{\text{rei}} = 10,000$ K.




FIG. 21.— Error in our code's calculation of $x_{\text{HeII}}(\Delta t)$ in units of the number of recombinations plotted as a function of Γ_{HeII} . This calculation assumes $\Delta t = 20$ Myr. The blue line represents the case $x_{\text{HeII}}(0) = 0.01$ and $\delta_b = 0$, the green line is for $x_{\text{HeII}}(0) = 1$ and $\delta_b = 0$, and the red line that for $x_{\text{HeII}}(0) = 1$ and $\delta_b = 10$. If we take n_e to be constant (a relatively good approximation during $z \sim 3$ HeII reionization because hydrogen is ionized), equation (C1) yields

$$x_{\text{HeII}}(t) = x_{\text{HeII,eq}} + (x_{\text{HeII}}(0) - x_{\text{HeII,eq}}) \exp[-t/t_{\text{eq}}], \quad (\text{C2})$$

where $t_{\text{eq}} = (\Gamma_{\text{HeII}} + 1/t_{\text{rec}})^{-1}$, $x_{\text{HeII}}(0)$ is the initial HeII fraction, and

$$x_{\text{HeII,eq}} = \frac{\alpha_A(T) n_e}{\Gamma_{\text{HeII}} + \alpha_A(T) n_e}. \quad (\text{C3})$$

Compare the exact solution for $x_{\text{HeII}}(t)$ given by equation (C2) to that without recombinations (and where an estimate for the number of recombinations $\alpha_A n_e x_{\text{HeII}}(0) \Delta t$ is added to x_{HeII} at the beginning of the timestep):

$$\tilde{x}_{\text{HeII}}(t) = [x_{\text{HeII}}(0) + \alpha_A n_e x_{\text{HeII}}(0)] \exp[-\Gamma_{\text{HeII}} t]. \quad (\text{C4})$$

The value $\tilde{x}_{\text{HeII}}(t)$ is what is computed by our HeII reionization code over a timestep Δt . The computation of $\tilde{x}_{\text{HeII}}(t)$ allows us to determine the effect of all sources on a single cell independently – Γ_{HeII} in equation (C4) is the sum of the Γ_{HeII} from all the incident rays – aside from the issue of the order rays from different sources reach a cell which affects the heating rate in the cell and how the cell shadows other cells.

Figure 21 shows that the error generated by using equation (C4) rather than the true solution, equation (C3), is minor. This figure plots the ratio of Δx_{HeII} to an estimate for the number of recombinations in a timestep $\alpha_A n_e \Delta t$ as a function of Γ_{HeII} for $\Delta t = 20$ Myr. The blue line is for $x_{\text{HeII}}(0) = 0.01$ and $\delta_b = 0$, the green line for $x_{\text{HeII}}(0) = 1$ and $\delta_b = 0$, and the red line for $x_{\text{HeII}}(0) = 1$ and $\delta_b = 10$. The error in Δx_{HeII} in all cases is small, less than 0.3 in units of the number of recombinations (and $\Delta x_{\text{HeII}} \lesssim 0.005$). For other values of Δt , we find that $\Delta x_{\text{HeII}}/(\alpha_A n_e \Delta t)$ peaks at roughly the same amplitude of 0.3, but with the peak at $\Gamma_{\text{HeII}} \approx 1.8 \Delta t^{-1}$.

This method also reproduces the correct heating rate with minimal error. It is easy to check this in the case of ionization equilibrium, the regime where one might suspect that this code is unreliable. In ionization equilibrium, the amount of heating is

$$\Delta Q = \Gamma_{\text{HeII}} \langle \Delta E \rangle x_{\text{HeII,eq}} \Delta t = \alpha_A n_e \langle \Delta E \rangle \Delta t, \quad (\text{C5})$$




FIG. 22.— Comparison of the dark matter and gas volume-averaged PDFs for $z = 3.5$. The solid curves are the dark matter PDFs and the dashed curves are the gas PDFs measured in the specified cell sizes. The dot-dashed curves are the MHR gas PDF.

where $\langle \Delta E \rangle$ is the average excess energy above 4 Ry of a photon that enters a cell and ionized a HeII ion, and we have used equation C3 for $x_{\text{HeII,eq}}$ in the limit $\Gamma_{\text{HeII}} \gg \alpha_A n_E$. Comparing ΔQ in equation C5 to the value our algorithm provides

$$\Delta Q = \Gamma_{\text{HeII}} \langle \Delta E \rangle \Delta x_{\text{HeII,rec}} = \alpha_A n_e \langle \Delta E \rangle \Delta t, \quad (\text{C6})$$

where we have used that $\Delta x_{\text{HeII,rec}} \approx \alpha_{A,\text{HeII}} n_e \Delta t$ (since $x_{\text{HeII}} \approx 1$) and that all the HeIII ions that recombine are reionized (i.e., $\Gamma_{\text{HeII}} \gg 1/t_{\text{rec}}$).

Since an accurate estimate for x_{HeII} is also necessary for HeII Ly α forest calculations, we correct the mistake our approximate method incurs by substituting the true solution, i.e. equation (C1), for cells that have $\tau_{\text{HeII,Ly}\alpha}(\tilde{x}_{\text{HeII}}, z, \delta_b) > 10$ (which correspond to larger values of $\tau_{\text{HeII,Ly}\alpha}(x_{\text{HeII}}, z, \delta_b)$), using the formula

$$\tau_{\text{HeII,Ly}\alpha}(\tilde{x}_{\text{HeII}}, z, \delta_b) = 3.4 \left(\frac{x_{\text{HeII}}}{10^{-3}} \right) \left(\frac{1+z}{4} \right)^{3/2} \Delta_b. \quad (\text{C7})$$

This correction might be worrisome because the algorithm no longer consistently calculates the balance between recombinations and ionizations. However, a minor correction of $\Delta x_{\text{HeII}} = |x_{\text{HeII}}(\Delta t) - \tilde{x}_{\text{HeII}}(\Delta t)|$ in a cell is acceptable because it is smaller than other uncertainties, such as uncertainties in the recombination coefficient (Section 2.5).

While we have used Case A recombination coefficients in the tests above, the equivalent discussion applies if we replace Case A with Case B, as is normally done by our code (except Δx_{HeII} will be smaller for Case B). We always use the Case A coefficient to correct our code's calculation because ground state recombination photons are not absorbed locally in regions with significant transmission. Rather, these photons are absorbed in dense (probably neutral) systems. As a consequence, our calculation of $\tau_{\text{HeII,Ly}\alpha}$ misses the contribution to Γ from ground state recombinations and, therefore, underestimates the amount of transmission. We can put an upper limit on the contribution to Γ_{HeII} from such recombinations. If we assume that the volume-averaged photo-ionization rate is just large enough to balance the number of recombinations aside from those to the ground state and that 55 eV recombination photons have the same m.f.p. as the typical ionizing photons from quasars, then $\bar{\Gamma}_{\text{rec,HeII}} = (\alpha_A - \alpha_B)/\alpha_A \bar{\Gamma}_{\text{HeII}} \approx 0.3 \bar{\Gamma}_{\text{HeII}}$. In our simulations, the photo-ionization rate is larger than what is required to balance the number of recombinations (except in select regions) and the m.f.p. for HeII-ionizing recombination photons is shorter than for other photons, such that $\bar{\Gamma}_{\text{rec,HeII}} < 0.3 \bar{\Gamma}_{\text{HeII}}$.

D. DARK MATTER VERSUS GAS DISTRIBUTION AND THE N_{HI} DISTRIBUTION

The radiative transfer computations in this study are performed as post processing on top of gridded N-body fields. This assumes that the gas fluctuations are pressure smoothed at roughly the simulation grid scale and trace the dark matter on larger scales. Cold dark matter clumps on extremely small scales, and, therefore, it is important to properly choose the scale at which to grid the N-body field to approximate the gas pressure smoothing. Here, we justify our grid sizes.

We use the Q6 simulation of Springel & Hernquist (2003). This simulation consists of a 14.3 Mpc box with 486^3 dark matter particles and initially 486^3 SPH particles. The Q6 simulation includes gas cooling, star formation, the Haardt & Madau (1996) model for the ionizing background (Katz et al. 1996a; Davé et al. 1999), as well as a prescription for galactic winds (Springel & Hernquist 2002). It resolves the Jeans scale at mean density with $\approx 10^4$ SPH particles, and it has a force resolution of 1.2 proper kpc. The following analysis uses the $z = 3.5$ snapshot from this simulation.

We first calculate the dark matter and gas overdensity PDFs to quantify how well the gridded dark matter traces the gas. This is shown in Figure 22 for 0.45 and 0.89 comoving Mpc grid cells. These cell sizes are comparable to the respective cell sizes for the 512^3 and 256^3 grids in the 186 Mpc simulation box. The solid curves are the dark




FIG. 23.— Number of absorbing columns per z per $\log_{10} N_{\text{HI}}$. The blue dotted curves represent the case in which this quantity is measured in 0.89 Mpc coarse cells. The thicker curve uses the average dark matter density in the coarse cell to measure N_{HI} , and the thin curve uses the gas density along a skewer the length of the coarse cell of the pressure-smoothed gas. The green dashed and red solid curves are the same as the blue dotted but measured in 0.45 and 0.22 Mpc coarse cells, respectively.

matter PDFs on the grid and the dotted are the same for the gas. The dot-dashed curve is the Miralda-Escudé et al. (2000) (henceforth MHR) gas PDF. This PDF represents that of the pressure smoothed gas (grid artifacts should be unimportant). The extent to which the gridded dark matter PDF agrees with the MHR gas PDF reflects how well the clumpiness of the gas is captured in the simulations.²⁷ The low density gas PDF is important for accurately calculating from the simulations the absorption in the HeII Ly α forest, the number of recombinations per HeII-ionizing photon, as well as the filtering of the radiation field. Figure 22 demonstrates that the gas clumpiness is described reasonably well with the chosen grid scales.

The high-density density regions, in particular densities responsible for absorbers with $N_{\text{HI}} \gtrsim 10^{15} \text{ cm}^{-2}$, are important for filtering the HeII radiation field (as shown in Appendix A). Figure 23 compares the N_{HI} column density distribution of the coarse cell-smoothed dark matter (thin curves) to the column density distribution of the pressure-smoothed gas (thick curves), with both column densities measured along skewers that have the same length as a coarse cell. Note that $N_{\text{HI}} = x_{\text{HI,eq}} l_{\text{cell}} \bar{n}_H(z) \Delta_b$, where we use $\Gamma_{\text{HI}} = 10^{-12} \text{ s}^{-1}$. The blue dotted curves represent the case in which N_{HI} is measured with 0.9 Mpc coarse cells. The green dashed and red solid curves are the same but measured in 0.45 and 0.22 Mpc coarse cells, respectively. The dark matter field, which for coarser grids underpredicts the number of $N_{\text{HI}} > 10^{15} \text{ cm}^{-2}$ absorbers, overproduces the abundance of absorbers for 0.22 Mpc cells. Interestingly, the N_{HI} distribution of the 0.45 Mpc gridded field (which corresponds roughly to the 512^3 grid simulation in the 186 Mpc box) agrees well with the gas distribution.

E. QSO MODELLING

QSO Spectral Index

Both methods for populating the box with QSOs discussed in Section 3 return the specific luminosity at the HI Lyman-limit. We extrapolate from this specific luminosity to higher energies using a power-law with index α_{UV} . The parameter α_{UV} has been measured by Telfer et al. (2002) to be ≈ 1.6 , but with an approximately Gaussian distribution with a large standard deviation (s.d.) of 0.86 (Telfer et al. 2002). The large s.d. in α_{UV} among QSOs may help explain transmission fluctuations in the HeII Ly α forest at $z \approx 2.5$ (Shull et al. 2004). More recently, Scott et al. (2004) measured a much smaller mean α_{UV} from a sample of both FUSE and HST quasars. They derived an average value of $\alpha_{\text{UV}} \approx 0.6$ and with a variance comparable to that of Telfer et al. (2002). Scott et al. (2004) also find a strong correlation between α_{UV} and QSO luminosity. Scott et al. (2004) argued that the discrepancy between their mean α_{UV} and that of Telfer et al. (2002) likely owed to their sample containing lower luminosity quasars.

The Telfer et al. (2002) and Scott et al. (2004) values of α_{UV} are derived from fitting $1 \text{ Ry} \lesssim E_\gamma \lesssim 4 \text{ Ry}$. For HeII reionization, the spectrum between 4 Ry and 1 keV is relevant. If we naively extrapolate from 1 Ry to the soft X-ray using the Telfer et al. (2002), or particularly the Scott et al. (2004), α_{UV} distribution this would result in an over-estimate of the soft X-ray luminosity function for quasars because of the large variance these studies measure in α_{UV} (and in the case of Scott et al. (2004), the low value of $\bar{\alpha}_{\text{UV}}$). The average spectral index $\alpha_{0,x}$ needed to join the luminosity at 2500 Å and 2 keV has been measured from a sample of optically selected quasars to be ~ 1.6 (Steffen et al. 2006).²⁸ Furthermore, Steffen et al. (2006) measured a dispersion in this index 0.08 – 0.14, significantly smaller than the dispersion in α_{UV} that Telfer et al. (2002) and Scott et al. (2004) infer.

²⁷ Note that the MHR PDF is generated using Eulerian simulations that employ a different cosmology and thermal history from simulation Q6. While the gas PDF is fairly robust to the cosmology, the corresponding PDF of the pressure smoothed gas in simulation G4 will be slightly different than the MHR PDF. We choose to use the MHR PDF because we find that particle noise in SPH simulations makes it difficult to capture the low density PDF.

²⁸ Steffen et al. (2006) found that this power-law was correlated with luminosity, ranging from $\alpha_{0,x} = 1.62$ for $L_{\text{bol}} = 10^{43} \text{ erg s}^{-1}$ and

Escape Fraction

Not all of the ionizing radiation escapes from a quasar. The fraction of HeII-ionizing radiation that is obscured is denoted by f_{cov} . Measurements in the X-ray suggest $f_{\text{cov}} \approx 0.5$, with evidence that f_{cov} increases with luminosity (Gilli et al. 2007). We use the Gilli et al. (2007) fitting function for f_{cov} in calculations that require it.²⁹

The fraction of ionizing radiation that escapes from QSOs is almost certainly directionally dependent – it will be easiest for ionizing photons to escape on trajectories that do not intersect infalling material or galactic disks. In select simulations, we take a beam for the ionizing radiation from the QSO with solid angle $\Omega_b = 2\pi(1 - f_{\text{cov}})$ along both axes. We assume the symmetry axis of the beam is randomly oriented. For the other simulations, we uniformly suppress the intensity along all sight-lines from the QSO by the factor f_{cov} . This approximates the scenario in which the obscuration owes to infalling material and (when averaged over the QSO lifetime) is isotropic. The value of the covering fraction is only required for Method I if the quasars are beamed. It affects the quasars in Method II by suppressing their ionizing luminosity.

REFERENCES

- Abel, T., & Haehnelt, M. G. 1999, *ApJL*, 520, L13
 Abel T., Wandelt B. D., 2002, *MNRAS*, 330, L53
 Agafonova I. I., Centuri  n M., Levshakov S. A., Molaro P., 2005, *A&A*, 441, 9
 Agafonova I. I., Levshakov S. A., Reimers D., Fechner C., Tytler D., Simcoe R. A., Songaila A., 2007, *A&A*, 461, 893
 Altay G., Croft R. A. C., Pelupessy I., 2008, *MNRAS*, pp 507+
 Anderson S. F., Hogan C. J., Williams B. F., Carswell R. F., 1999, *AJ*, 117, 56
 Bajtlik, S., Duncan, R. C., & Ostriker, J. P. 1988, *ApJ*, 327, 570
 Becker G. D., Rauch M., Sargent W. L. W., 2007, *ApJ*, 662, 72
 Bernardi M., Sheth R. K., SubbaRao M., Richards G. T., Burles S., Connolly A. J., Frieman J., Nichol R., Schaye J., Schneider D. P., Vanden Berk D. E., York D. G., Brinkmann J., Lamb D. Q., 2003, *AJ*, 125, 32
 Boksenberg A., Sargent W. L. W., Rauch M., 2003, Submitted to *Astrophysical Journal Supplement*. astro-ph/0307557
 Bolton, J. S., Haehnelt, M. G., Viel, M., & Springel, V. 2005, *MNRAS*, 357, 1178
 Bolton J. S., Haehnelt M. G., Viel M., Carswell R. F., 2006, *MNRAS*, 366, 1378
 Bolton, J. S., Viel, M., Kim, T.-S., Haehnelt, M. G., & Carswell, R. F. 2008a, *MNRAS*, 386, 1131
 Bolton, J. S., Oh, S. P., & Furlanetto, S. R. 2008b, arXiv:0807.2447
 Bouwens, R. J., Illingworth, G. D., Franx, M., & Ford, H. 2007, *ApJ*, 670, 928
 Cen R., Miralda-Escud   J., Ostriker J. P., Rauch M., 1994, *ApJL*, 437, L9
 Cen R., Simcoe R. A., 1997, *ApJ*, 483, 8
 Conroy, C., Wechsler, R. H., & Kravtsov, A. V. 2006, *ApJ*, 647, 201
 Croft R. A. C., 2004, *ApJ*, 610, 642
 Croom S. M., Boyle B. J., Shanks T., Smith R. J., Miller L., Outram P. J., Loaring N. S., Hoyle F., da Angela J., 2005, *MNRAS*, 356, 415
 Dav  , R., Hernquist L., Katz N., Weinberg D. H. 1999, *ApJ*, 511, 521
 Daviden A. F., Kriss G. A., Zheng W., 1996, *Nature*, 380, 47
 Fan X., Narayanan V. K., Strauss M. A., White R. L., Becker R. H., Pentericci L., Rix H.-W., 2002, *AJ*, 123, 1247
 Fardal M. A., Giroux M. L., Shull J. M., 1998, *AJ*, 115, 2206
 Faucher-Gigu  re C.-A., Prochaska J. X., Lidz A., Hernquist L., Zaldarriaga M., 2008a, *ApJ*, 681, 831
 Faucher-Gigu  re C.-A., Lidz A., Hernquist L., Zaldarriaga M., 2008c, *ApJ*, in prep.
 Fechner C., et al., 2006, *A&A*, 455, 91
 Furlanetto S., Oh S. P., 2007, *ArXiv e-prints*, 711
 Furlanetto, S. R., & Oh, S. P. 2008, *ApJ*, 681, 1
 Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, *ApJ*, 613, 1
 Gilli R., Comastri A., Hasinger G., 2007, *A&A*, 463, 79
 Giroux M. L., Shull J. M., 1997, *AJ*, 113, 1505
 Gleser L., Nusser A., Benson A. J., Ohno H., Sugiyama N., 2005, *MNRAS*, 361, 1399
 Gnedin N. Y., Hui L., 1998, *MNRAS*, 296, 44
 Gnedin N. Y., Kravtsov A. V., Chen H.-W., 2008, *ApJ*, 672, 765
 Haardt F., Madau P., 1996, *ApJ*, 461, 20
 Heap S. R., Williger G. M., Smette A., Hubeny I., Sahu M., Jenkins E. B., Tripp T. M., Winkler J. N., 2000, *ApJ*, 534, 69
 Hernquist L., Katz N., Weinberg D. H., Miralda-Escud   J., 1996, *ApJL*, 457, L51+
 Hogan C. J., Anderson S. F., Rugers M. H., 1997, *AJ*, 113, 1495
 Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Martini P., Robertson B., Springel V., 2005, *ApJ*, 630, 705
 Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2006, *ApJS*, 163, 1
 Hopkins P. F., Richards G. T., Hernquist L., 2007, *ApJ*, 654, 731
 Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Krause E., 2007a, *ApJ*, 669, 45
 Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Krause E., 2007b, *ApJ*, 669, 67
 Hopkins P. F., Hernquist L., Cox T. J., Keres, D. 2008a, *ApJS*, 175, 356
 Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
 Hui L., Haiman Z., 2003, *ApJ*, 596, 9
 Jakobsen P., Boksenberg A., Deharveng J. M., Greenfield P., Jedrzejewski R., Paresce F., 1994, *Nature*, 370, 35
 Kashikawa N., et al., 2006, *ApJ*, 648, 7
 Katz N., Weinberg D. H., Hernquist L., Miralda-Escud   J., 1996, *ApJL*, 457, L57
 Katz N., Weinberg D. H., Hernquist L. 1996a, *ApJS*, 105, 19
 Komatsu E., Dunkley J., Nolte M. R., Bennett C. L., Gold B., Hinshaw G., Jarosik N., Larson D., Limon M., Page L., Spergel D. N., Halpern M., Hill R. S., Kogut A., Meyer S. S., Tucker G. S., Weiland J. L., Wollack E., Wright E. L., 2008, *ArXiv e-prints*, 803
 Lai K., Lidz A., Hernquist L., Zaldarriaga M., 2006, *ApJ*, 644, 61
 Lidz A., McQuinn M., Zaldarriaga M., Hernquist L., Dutta S., 2007, *ApJ*, 670, 39
 Madau P., Haardt F., Rees M. J., 1999, *ApJ*, 514, 648
 Madau P., Rees M. J., Volonteri M., Haardt F., Oh S. P., 2004, *ApJ*, 604, 484
 McDonald P., Miralda-Escud   J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, *ApJ*, 562, 52
 McDonald P., Seljak U., Cen R., Bode P., Ostriker J. P., 2005, *MNRAS*, 360, 1471
 McQuinn M., Lidz A., Zahn O., Dutta S., Hernquist L., Zaldarriaga M., 2007, *MNRAS*, 377, 1043
 Meiksin A., White M., 2004, *MNRAS*, 350, 1107
 Miniati F., Ferrara A., White S. D. M., Bianchi S., 2004, *MNRAS*, 348, 964
 Miralda-Escud   J., & Rees, M. J. 1994, *MNRAS*, 266, 343
 Miralda-Escud   J., Haehnelt M., Rees M. J., 2000, *ApJ*, 530, 1
 Miralda-Escud   J., 1998, *ApJ*, 501, 15
 Miralda-Escud   J., Cen R., Ostriker J. P., Rauch M., 1996, *ApJ*, 471, 582
 Paschos P., Norman M. L., Bordner J. O., Harkness R., 2007, *ArXiv e-prints*, 711
 Pawlik A. H., Schaye J., 2008, *ArXiv e-prints*, 802
 Petitjean P., Webb J. K., Rauch M., Carswell R. F., Lanzetta K., 1993, *MNRAS*, 262, 499
 Porciani C., Magliocchetti M., Norberg P., 2004, *MNRAS*, 355, 1010

$\alpha_{0,x} = 1.25$ for $L_{\text{bol}} = 10^{46}$ erg s⁻¹. This trend with L_{bol} is the opposite of what was found for α_{UV} in Scott et al. (2004).

²⁹ However, f_{cov} is defined in Gilli et al. (2007) as the fraction of QSOs with $\log_{10}(N_{\text{HI}}) > 21$ in c.g.s. units, which is not exactly the fraction of radiation that escapes, which for is sensitive to the number of sight-lines with $\log_{10}(N_{\text{HeII}}) > 17.9$.

- Reimers D., Fechner C., Hagen H.-J., Jakobsen P., Tytler D., Kirkman D., 2005, *A&A*, 442, 63
- Reimers D., Kohler S., Wisotzki L., Groote D., Rodriguez-Pascual P., Wamsteker W., 1997, *A&A*, 327, 890
- Ricotti M., Gnedin N. Y., Shull J. M., 2000, *ApJ*, 534, 41
- Schaye J., 2001, *ApJ*, 559, 507
- Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
- Scott, J., Bechtold, J., Dobrzycki, A., & Kulkarni, V. P. 2000, *ApJS*, 130, 67
- Scott J. E., Kriss G. A., Brotherton M., Green R. F., Hutchings J., Shull J. M., Zheng W., 2004, *ApJ*, 615, 135
- Shapley A. E., Steidel C. C., Pettini M., Adelberger K. L., 2003, *ApJ*, 588, 65
- Shull J. M., Tumlinson J., Giroux M. L., Kriss G. A., Reimers D., 2004, *ApJ*, 600, 570
- Shull J. M., van Steenberg M. E., 1985, *ApJ*, 298, 268
- Sokasian A., Abel T., Hernquist L., 2001, *NewA*, 6, 359
- Sokasian A., Abel T., Hernquist L., 2002, *MNRAS*, 332, 601
- Sokasian A., Abel T., Hernquist L., 2003, *MNRAS*, 340, 473
- Songaila A., 1998, *AJ*, 115, 2184
- Springel V., Hernquist L., 2002, *MNRAS*, 333, 649
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 312
- Stark, D. P., Ellis, R. S., Richard, J., Kneib, J.-P., Smith, G. P., & Santos, M. R. 2007, *ApJ*, 663, 10
- Steffen A. T., Strateva I., Brandt W. N., Alexander D. M., Koekemoer A. M., Lehmer B. D., Schneider D. P., Vignali C., 2006, *AJ*, 131, 2826
- Steidel C. C., Pettini M., Adelberger K. L., 2001, *Astrophys. J.*, 546, 665
- Taniguchi Y., et al., 2005, *PASJ*, 57, 165
- Telfer R. C., Zheng W., Kriss G. A., Davidsen A. F., 2002, *ApJ*, 565, 773
- Theuns, T., Zaroubi, S., Kim, T.-S., Tzanavaris, P., & Carswell, R. F. 2002a, *MNRAS*, 332, 367
- Theuns T., Bernardi M., Frieman J., Hewett P., Schaye J., Sheth R. K., SubbaRao M., 2002b, *ApJL*, 574, L111
- Tinker, J. L., Weinberg, D. H., Zheng, Z., & Zehavi, I. 2005, *ApJ*, 631, 41
- Tittley, E. R., & Meiksin, A. 2007, *MNRAS*, 380, 1369
- Totani T., Kawai N., Kosugi G., Aoki K., Yamada T., Iye M., Ohta K., Hattori T., 2006, *PASJ*, 58, 485
- Trac H., Cen R., 2007, *ApJ*, 671, 1
- Vale, A., & Ostriker, J. P. 2006, *MNRAS*, 371, 1173
- Zahn, O., Lidz, A., McQuinn, M., Dutta, S., Hernquist, L., Zaldarriaga, M., & Furlanetto, S. R. 2007, *ApJ*, 654, 12
- Zaldarriaga, M. 2002, *ApJ*, 564, 153
- Zaldarriaga M., Hui L., Tegmark M., 2001, *ApJ*, 557, 519
- Zheng W., Chiu K., Anderson S. F., Schneider D. P., Hogan C. J., York D. G., Burles S., Brinkmann J., 2004, *AJ*, 127, 656
- Zuo L., Phinney E. S., 1993, *ApJ*, 418, 28